



Confidence in a connected world.

Managing Storage Complexity for Healthcare Providers

Managing Storage Complexity for Healthcare Providers

Contents

Executive summary	4
Introduction: Valuable technology, vast amounts of information	5
EMR and EHR challenges	5
Improving Storage Resource Management (SRM)	7
Improving utilization of existing storage	7
Massive storage requirements result in massive backups	8
Eliminating duplicated data for faster backups	8
Protecting stored data	9
Adopting an integrated solution to meet the challenge	9
Providing access and availability while controlling IT costs	10
Prioritizing data and using tiered storage to optimize backups	11
Ensuring reliable, efficient data backup and recovery	13
Anticipating future needs with scalable systems	14
Highlights of Symantec solutions	15
Conclusion	16

Executive summary

This paper is one of a series of papers that focus on the challenges healthcare providers face in information technology and the best practices for meeting those challenges. This paper focuses on the specific topic of data storage management.

Advances in diagnostic technology and the increasing popularity of digital medical images and records have resulted in large networks of heterogeneous storage systems that now house vast amounts of stored data. With storage requirements increasing exponentially every year, how can providers take a proactive approach to meeting future data storage challenges?

This white paper discusses an effective data management strategy that will allow healthcare providers to maximize storage utilization and realize the full benefits of digital technology while keeping IT costs under control. With an enterprise-wide strategy and implementation plan for records management and medical image archiving, IT managers can use a single interface to view, control, and migrate information stored on multiple servers in disparate locations. A centralized management solution also simplifies procedures for archiving, backing up, and recovering data according to organizational and regulatory requirements. By prioritizing information and creating different tiers of storage, administrators can place low-priority data on less costly storage while ensuring that recent information is easily retrievable. Eliminating redundant backups of unique data can reduce storage requirements and shorten backup window times, thereby lessening demands on IT and lowering costs.

A centralized enterprise architecture for storage improves utilization, relieves demands on IT, and lowers costs. Although clinical imaging data alone is a large part of the problem, a comprehensive solution incorporates best practices for backup and storage of medical images, electronic medical records (EMRs), and all types of digital information across the healthcare industry.

Introduction: Valuable technology, vast amounts of information

When advances in digital medical imaging and EMR technologies began to change the way hospitals created and managed patient healthcare information, the medical community quickly realized how the technology would improve diagnosis and treatment for patients. What wasn't so readily apparent was the challenge that digitization would pose for IT departments in managing the resultant mountains of data.

Early adoption of medical imaging began in radiology departments, where hospitals deployed stand-alone systems and dedicated hardware solutions to manage storage of digitized x-rays, MRIs, CAT scans, and other radiological images. As the benefits of digitizing became obvious and new technologies emerged, the popularity of medical imaging spread to other clinical areas such as cardiology, pathology, and dermatology, thus driving demand for even more deployments of multiple storage islands within a single institution.

Although storage systems were successful within individual departments, they filled to capacity quickly. Before long, IT departments found themselves managing multiple systems on large networks of heterogeneous storage systems that were not scalable or integrated. Application servers and storage islands have been useful to a certain extent, but they are highly specialized and they cannot communicate with each other, work with each other, or share resources.

EMR and EHR challenges

Federal mandates in recent years have focused on digitizing health records and other medical information to reduce errors, lower costs, improve patient care, and help facilitate clinical data exchange. As national agendas have spurred the implementation of state and local incentives and digital technology has become more sophisticated, the transition of hard-copy patient information to EMRs has already begun in larger and mid-sized institutions (see Figure 1). Analysts anticipate widespread adoption of EMRs within the next decade.¹

¹ IDC Health Industry Insights US HIT Market Overview, Marc Holland, Program Director, Healthcare Provider Research, May 17, 2007.

Managing Storage Complexity for Healthcare Providers

Hospitals adopted imaging technology much more quickly and easily than EMRs, but storage challenges have accompanied the obvious benefits of EMRs. Although single provider organizations “own” EMRs, they are frequently used, annotated, and stored in multiple locations. EMRs are often included in patients’ electronic health records (EHRs), which contain clinical information from multiple provider organizations (see Figure 1).

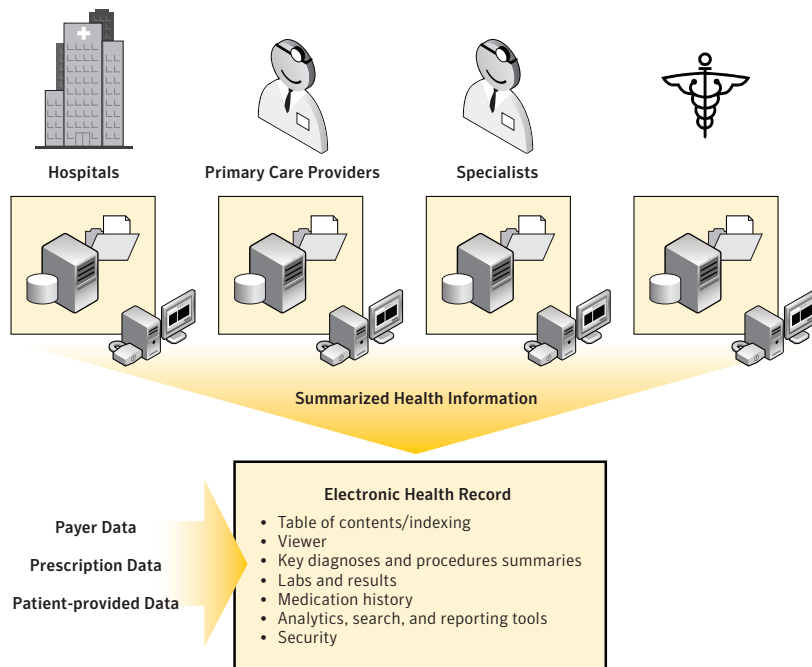


Figure 1. EMRs and EHRs contain digitized images from many sources

Source: IDC Health Industry Insights.²

EMRs and EHRs summarize and organize patient information, including digitized images of paper documents and electronic data from patients, payers, and pharmacies, all of which can form the basis of a longitudinal medical record.³ Both EMRs and EHRs contain vast amounts of redundant form-based information (structured data) that must be copied into backup and disaster recovery versions.

² Ibid.
³ Ibid.

Improving Storage Resource Management (SRM)

Storage software is often packaged with the various systems that healthcare organizations purchase from different application vendors, and each system requires independent maintenance and administration. In many cases, separate storage systems are controlled by different operating systems—Windows®, UNIX, Linux®, AIX®, Sun™ Solaris™, HP-UX, and others—and the management capabilities of each storage island function only within that island's particular operating system. The result is that a hospital may manage storage in a variety of ways at the server level using different storage hardware vendors and different backup solutions, which can be costly and inefficient. For many hospitals, storage demands are growing at up to 70 percent each year,⁴ and current data storage systems are not scalable to meet the demands of exponentially increasing amounts of retained data.

With the spread of advanced medical imaging technologies to other departments and the growing demand for high-quality healthcare, storage requirements have placed heavy demands on IT and network resources. A single x-ray can take up to 12 megabytes of storage, and a 500-bed hospital can easily generate 1 to 2 terabytes of storage per year in x-rays alone, with backed-up data doubling that requirement.⁵ The rapidly growing storage in hospitals translates to more IT staff resources needed to manage it, and the demand is especially burdensome due to different storage systems that have to be managed individually. Without an enterprise-wide storage solution, providers are continuing to purchase and deploy additional storage islands—each of which requires even more individual management. Implementation of a solution that centralizes the management of stored data using a single interface would maximize the utilization of these various storage systems to accommodate growing amounts of data, thereby reducing costs for purchasing additional storage hardware and relieving demands on IT.

Improving utilization of existing storage

Many providers are working with literally thousands of possible combinations of hardware and software in heterogeneous data centers, and storage space is not utilized efficiently. For example, medical data for one patient may be stored in separate servers in radiology, surgery, and medical records departments, and the data may be stored in different formats such as EMRs, x-rays, lab results, and scanned records. Each storage island usually contains only one type of data. If no centralized storage management solution is in place to optimize use of storage space across these islands, as much as 33 percent of available storage in each location may go unused.⁶

⁴ Raths, David. "Information Explosion." *Healthcare Informatics*, February 2007.

⁵ Moore, Fred. "Storage Intensive Applications—Industry Trend or Event." *Computer Technology Review*, April 2001.

⁶ Personal communication, Martin Ward, Symantec, October 5, 2007.

Managing Storage Complexity for Healthcare Providers

Another challenge is that hospitals want the ability to locate data related to a single patient and archive the related data in the same place for fast and easy access rather than having to retrieve each piece of data separately from the disparate location where it is stored. A centralized virtual storage management solution would provide visibility and control from one interface, allowing IT managers to prioritize files and create a plan for accessing and retaining data. With a comprehensive records management plan in place, administrators can consolidate data from disparate storage volumes into logical storage tiers to increase utilization and optimize accessibility of data.

Massive storage requirements result in massive backups

Because stored data must be backed up regularly, the exponential growth in storage requirements translates to challenging demands on traditional backup methods. According to industry estimates, storage capacity requirements for medical imaging data double every 18 months, as PACS is used outside the radiology department and the technology becomes more sophisticated.⁷ Large amounts of unstructured (anecdotal) data not specifically tied to a patient's protected health information (PHI) must nonetheless be stored and backed up. Emails and image file attachments also constitute huge amounts of redundant information that is stored and backed up.

Existing backup methods are not optimized, security procedures are not reliable, and overnight backup windows are insufficient due to the large amounts of data. Replication and existing online backup approaches designed to avoid backup windows place heavy demands on system resources that interfere with regular operations, but these problems can be reduced or entirely eliminated by a storage solution that provides for de-duplication of redundant data.

Eliminating duplicated data for faster backups

A major underlying cause of rapidly rising storage requirements is the amount of redundant data stored in multiple locations. Structured, form-based data is often repeated throughout numerous different departments, with identical patient information stored and backed up in storage islands in clinical data repositories (CDRs), medical records data banks, email servers, and other locations. In some facilities, duplicate data is backed up 50 to 500 times using common backup techniques.⁸ If hospitals can identify unique data and prevent it from being duplicated repeatedly, storage utilization will rise dramatically and backup times can be reduced.

⁷ Digital Imaging: Reducing Medical Error and Improving Diagnostics (Review Report), Datamonitor, February 2007.

⁸ Internal Symantec initiative, "Storage United."

Protecting stored data

One of the key advantages of digitization is that it enables hospitals to create many more images than they could with previously used imaging technologies. Again, the result of this benefit is massive amounts of data that must be stored and archived. Federal regulation standards, including the Health Insurance Portability and Accountability Act (HIPAA), require organizations to securely store healthcare records and develop comprehensive plans for backup, disaster recovery, and emergency medical operations to ensure that information is protected against loss. Storage protection must account for physical space allocations, power requirements, and temperature control systems.

In many cases, hospitals are required to store digital records for decades. For example, records for oncology patients never expire, and medical records for pediatric patients are usually stored at least until the patient turns 18 and often even longer. Because hospitals typically must retain digitized data for a longer period than the storage hardware's useful life, hospitals need the ability to migrate storage of older records to new systems while minimizing the risk of information damage or loss. An integrated solution can provide just such portability by allowing IT managers to move data to newly acquired storage hardware, regardless of the platform it currently exists on, using a single interface that can access and manage information across a heterogeneous storage network. A comprehensive information lifecycle management (ILM) plan should be implemented for classifying, protecting, and archiving data according to organizational retention requirements.

Adopting an integrated solution to meet the challenge

The key to conquering the myriad storage management problems that exist in the healthcare industry is understanding and addressing the root causes. A united storage solution for healthcare enterprises enables them to provide data storage as a valued service to patients, providers, and payers alike. Such a solution enables hospitals to manage not only the operational processes but also the technology behind them. By creating an integrated software infrastructure that is standardized, healthcare organizations can select the most appropriate operating system, disk subsystem, tape library, and applications, and then use the same software to manage all of the hardware components—no matter what vendor provided them—and access data from multiple locations (see Figure 2).

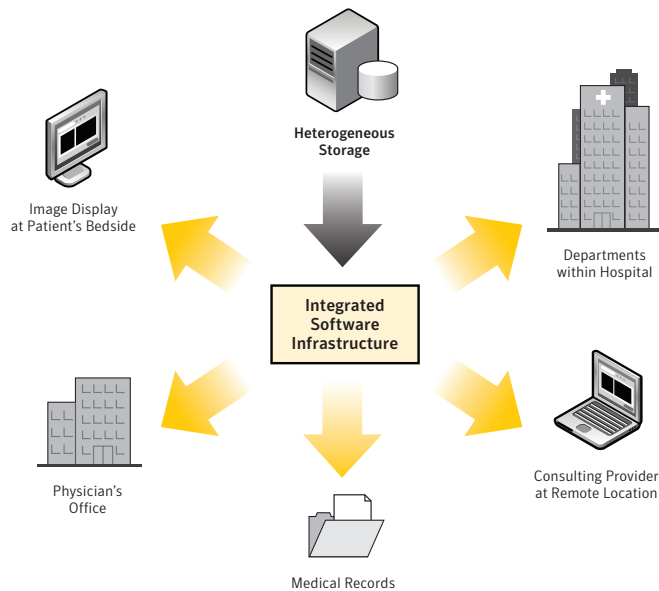


Figure 2. Integrated infrastructure solution.

Using one software interface across multiple platforms means that IT uses a centralized management console with one interface for all platforms. Storage managers can view different departments—radiology, cardiology, and others—and make storage decisions based on a high-level virtualized view of the entire storage network. And as servers, storage, and other hardware components are retired, they can be upgraded seamlessly without purchasing new software.

Providing access and availability while controlling IT costs

Ensuring fast access to stored data in a hospital can be a matter of life and death. Remote sites such as hospitals, clinics, and doctors' offices must be able to quickly and efficiently search for and access a complete set of digitized medical images, EMRs, and other stored information without involving IT personnel, using a familiar user interface. An effective storage solution for medical images and EMRs should give hospitals faster access to healthcare records from any location at any time, should allow data to be transferred seamlessly among remote locations without requiring IT involvement, and should provide a robust disaster recovery solution.

Managing Storage Complexity for Healthcare Providers

Children's Healthcare of Atlanta, a world-renowned pediatric hospital system serving the Atlanta metropolitan area and the entire state of Georgia, required a software infrastructure that would bring together multiple server platforms, operating systems, and storage systems to handle a rapid growth of data resulting from regulatory compliance requirements. The hospital was supporting three separate clustered server environments, and the lack of standardization required separate IT management of each of the software solutions.

High availability and data protection are top priorities at Children's. Implementing a standardized disaster recovery solution was critical, because the hospital wanted to avoid having to train specialized staff to manage each of the different storage areas. The IT team chose to deploy a centralized software infrastructure to manage the heterogeneous hardware environments, including a cost-effective data replication and application failover software solution.

This solution allowed the IT team to increase the utilization of existing storage so the hospital would not have to purchase additional hardware. Children's realized an immediate savings of \$2.3 million and a four-year total savings of \$4.7 million. Due to the deployment of hardware-independent disaster recovery software, ongoing storage costs were reduced by 53 percent. "One of the major benefits of the solution is standardization; any time you can get to standardization, there is going to be some savings there," said Jim Atwood, Director of Technology at Children's. "We can install on one technology, train on one technology, and know that we can support a variety of systems with that same technology."⁹

A standardized infrastructure dramatically decreases capital expenditures because hospitals can still maintain a multi-vendor environment and seek the best pricing for hardware, while reducing IT costs by centrally administering storage resources using the same skill set.

Prioritizing data and using tiered storage to optimize backups

Before IT managers can create storage policies, they need to know what information will be stored and how it will be used. Hospitals need fast access to recent diagnostic images, data required for surgeries or consultations, updated patient EMRs, and emails related to a patient's current treatment.

⁹ "Children's Healthcare of Atlanta: Measuring Business Value from High Availability and Disaster Recovery," E-Business Strategies, Inc. Business Value Analysis Study, 2005.

Managing Storage Complexity for Healthcare Providers

An effective storage solution employs dynamic storage tiering (DST) to place information on predefined storage tiers. Hospitals can save money by creating tiers of digitized medical images, EMRs, and other digitized data, using information lifecycle management (ILM) policies to prioritize data according to its value. IT managers can then determine what types and sizes of storage tiers are needed, and then design tiers according to the recovery time objectives (RTO) for each type of data (see Figure 3).

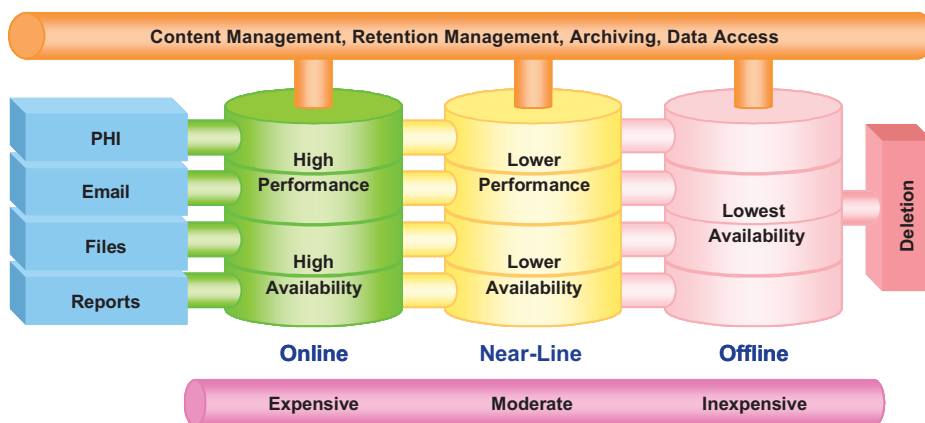


Figure 3. ILM migrates data to lower cost storage.¹⁰

Administrators can move files among tiers based on end-user policies or data retention policies, placing active patient data on fast-access storage for immediate retrieval. Information needed periodically, such as mammograms that are reviewed yearly, can be placed in mid-level storage tiers where files can be accessed easily. Inactive patient information, such as records for a deceased person, can be archived on low-cost storage. DST can benefit administrators in implementing an email management and archiving system to reduce storage requirements, which reduces both expense and complexity in maintaining storage systems. (See “Managing Electronic Messaging and E-Discovery for Healthcare Providers” in this Best Practices Series for a detailed discussion of email management.)

¹⁰Gartner, Inc., “Data Classification Is a Vital First Step in Information Life Cycle Management,” by Barry Runyon, June 27, 2007.

Managing Storage Complexity for Healthcare Providers

A well-designed DST strategy for prioritization and clustered storage of data gives healthcare organizations fast and reliable access to information as it is needed, and reduces the need for repeated backup of older data that is infrequently accessed. An optimal backup process moves information to the correct tier automatically and removes data as it is needed. Structured (clinical) data is already stored in specific tiers in most facilities, in order to comply with HIPAA regulations. But unstructured data, which accounts for up to 80 percent of stored data, should be identified and moved from tier to tier, stored in a data repository that feeds into the other storage areas.

In today's hospital environment, new data can be prioritized according to "value of data" policies the user provides as the digital data is created. But older data that existed before digital imaging does not contain such policies. In order to determine data retention policies for archived data, hospitals need to evaluate data according to its file attributes—the department that created it, the age of the data, when it was created, etc.—and then store it in the appropriate tier for later access or purge it from storage if it is no longer needed based on clinical, research, and/or legal retention requirements.¹¹

Community Health Network, a leading-edge healthcare provider consisting of five major regional hospitals in the Indianapolis area, was challenged to improve its data recovery and backup processes to ensure minimal downtime of critical applications. To meet this challenge, the IT department created three tiers of application criticality, each with specific data protection policies and recovery objectives. Tier 1 data and applications are synchronously replicated to the disaster recovery site in real time; tier 2 houses applications that need to be restored in a 24- to 48-hour period; and tier 3 applications are less critical and can be restored in a 48- to 72-hour period. Community Health Network's strategy of creating application tiers with criticality and recovery objectives provided a comprehensive and effective approach to data protection and recovery, ensuring 24x7 access to medical, clinical, and administrative information.¹²

Ensuring reliable, efficient data backup and recovery

A recommended solution will securely and centrally control, manage, and back up data residing on heterogeneous storage systems in disparate locations that have limited resources. A united storage platform supports backup of all types of data regardless of the hardware housing it. Consolidating the management of data backup under a single IT resource using a standardized software solution reduces costs and IT administrative time.

¹¹ Personal communication, Martin Ward, Symantec.

¹² DeBois, Laura, "Community Health Network: Delivering State-of-the-Art Patient Care with a Digitally Integrated Network." IDC Buyer Case Study, April 2007.

Managing Storage Complexity for Healthcare Providers

Advanced backup techniques that consolidate the duplication of structured data (de-duplication) greatly reduce the volume of stored healthcare information, particularly with structured, form-based data such as EMRs that have traditionally been copied multiple times into backup and disaster recovery versions. A backup solution that focuses on de-duplication of data also reduces the time required for backups so they can be completed in an overnight backup window. The ability to identify unique data that has already been backed up, both locally and remotely, eliminates redundancy, lowers hardware costs, reduces network consumption, and streamlines storage management and administration.

St. Mary's Regional Medical Center, a 233-bed acute care facility based in Lewiston, Maine, was conducting nightly backups using tape libraries on five backup servers, and they were quickly running out of space and time. With tape reliability and handling issues and a backup window of well over 12 hours, data retention and restoration was costly and resource-intensive. After deployment of a disk-based backup system implementing byte-level data de-duplication, St. Mary's achieved a 50:1 reduction in data storage requirements, reduced backup management from two to three hours per day to zero, and cut their backup window to just a few hours.¹³

Anticipating future needs with scalable systems

Robust, high-availability networks are one key to developing an effective solution to the mounting problems caused by rising storage requirements. As networking technology continues to evolve, hospitals will consider partial upgrading and later, complete new network installation.

As the environment grows in terms of remote offices (clinics, physicians' offices, etc.), physical space requirements, and common file servers, the hardware storage system should be scaled to accommodate growth without impacting storage resource management and administration. If all available storage is virtualized in logical storage pools across an integrated software infrastructure, scaling can occur online without reconfiguring backup and replication policies or clients. Hardware components can be replaced or upgraded seamlessly without affecting storage resource management (SRM) operations. And because the centralized software infrastructure increases utilization of existing storage, hospitals will not need to purchase new storage system hardware as often.

¹³ "St. Mary's Regional Medical Center: Disk Heals Backup Woes," Joe Longtin, St. Mary's Medical Center; Fred Pinkett, ExaGrid Systems; Mike Garcia, Symantec.

Managing Storage Complexity for Healthcare Providers

Highlights of Symantec solutions

The following table lists Symantec solutions for managing data storage.

Storage Resource Management Need	Symantec Solution
<ul style="list-style-type: none"> • Prioritize information and create data retention policies to determine where it should be stored; store data according to its RTO • Dynamically migrate data to different tiers of storage and perform seamless data migrations across different server architectures • Store recent files on fast, near-term storage and older files on less expensive systems • Monitor, manage, and report on data stored on different platforms using a single Web-based console • Reduce IT training and labor costs • Transfer data among remote locations without IT involvement • Reduce complexity, ensure availability, and improve efficiency and productivity without adding significant cost 	Veritas Storage Foundation™ Veritas™ Volume Manager Veritas™ File System
<ul style="list-style-type: none"> • Manage heterogeneous storage resources from a centralized location • Choose the appropriate storage hardware without compromising how it is managed • Improve utilization and control costs 	Veritas CommandCentral™ Storage
<ul style="list-style-type: none"> • Manage, protect, and recover data across different storage tiers, locations, and operating systems using a single interface • Meet backup windows with faster disk-based backups • Improve storage utilization • Centralize management and backup of heterogeneous storage systems in disparate locations using a single interface • Reduce IT management and administration time, and control costs 	Veritas NetBackup™
<ul style="list-style-type: none"> • Control and manage the retention and recovery of backup data from multiple locations for different platforms and applications using a standard file system interface • Reduce volume of structured, form-based data to eliminate redundant backup processes • Reduce use of network and system resources for backups • Reduce IT training and labor costs • Improve ROI with a scalable and open software-based storage system • Centralize data protection administration, management, and compliance by providing a reliable and consistent backup and recovery process • Comply with HIPAA regulations 	Veritas NetBackup PureDisk™
<ul style="list-style-type: none"> • Gain visibility into unused storage capacity and reclaim space to maximize the utilization of existing storage assets and defer expensive hardware purchases 	Consulting Services for Symantec Storage Optimization
<ul style="list-style-type: none"> • Implement proven storage management processes and policies and introduce automation to achieve operational efficiency and cost savings 	Consulting Services for Symantec Storage Process Transformation

Conclusion

Administering, storing, and managing vast amounts of healthcare data is an ongoing challenge, and the demands continue to grow. Symantec suggests that hospitals adopt a proactive strategy for managing stored information efficiently to control costs and realize the full benefits offered by digital imaging, EMRs, and other technological advances. A best practices approach should include the following objectives:

- Develop an enterprise-wide strategy and implementation plan for standardization of records management, email storage, and medical image archiving.
- Centralize storage management to improve utilization, relieve demands on IT, and reduce costs.
- Create procedures for discovering, identifying, and prioritizing critical enterprise data, both structured and unstructured.
- Implement de-duplication processes to reduce the amount of redundant data stored in multiple locations.
- Develop a tiered storage infrastructure as preparation for limited or full-scale ILM; the various types of data found within the enterprise and associated service-level objectives will determine the number of tiers required.
- Implement a centralized backup infrastructure for secure protection of data on heterogeneous storage systems.
- Establish formal data backup and retention schedules that align the business value of classified information with storage tiers.

By implementing an integrated, standardized storage infrastructure, healthcare organizations can achieve these objectives and thereby streamline patient care, reduce medical error, control costs, and increase efficiency.

About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world.

The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries.

More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

Copyright © 2007 Symantec Corporation. All rights reserved. Symantec, the Symantec Logo, CommandCentral, NetBackup, PureDisk, Veritas, and Veritas Storage Foundation are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. AIX is a trademark of International Business Machines Corporation in the United States, other countries, or both. Windows is either a registered trademark or a trademark of Microsoft Corporation in the United States and/or other countries. Sun and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States or other countries. Other names may be trademarks of their respective owners. Printed in the U.S.A. 11/07 13524275