

Multi-Terabyte Archives for Medical Imaging Applications

This paper describes how Windows servers running XenData Archive Series software provide an attractive solution for storing and retrieving multiple terabytes of medical image files. The XenData solution saves data to both RAID and tape, taking advantage of the characteristics of each media type. The solution overcomes the limitations of conventional backup and optical disk based systems.

Copyright 2005 XenData Limited. All rights reserved.

Introduction

XenData Archive Series software delivers a unique solution for archiving multiple terabytes of data. It manages RAID and a tape library on a Windows server, combining hierarchical storage management, integrated data protection and file version management. It is optimized for archiving and retrieving large volumes of data, scaling from terabytes to over a petabyte per server and is ideal for archiving medical images.

Used alone, XenData software works at the file level and presents the tape library and RAID as a standard Windows file system. When interfaced with a DICOM server, the solution will adhere to the network transmission and management protocol for DICOM images and will store data in DICOM format¹.

Limitations of Conventional Backup

For small medical image archives, storing data to magnetic disk and backing up to tape for disaster recovery purposes works well. However, as the volume of stored data increases this approach becomes more and more difficult to implement. The major problem is the time to complete a full backup. For example, it takes 13.8 hours to perform a full backup of 1TB of data at an average transfer rate of 20 megabytes per second (MB/s) and over 5 days to complete a full backup of 10 TB. The backup window problem can be alleviated to some extent by employing more tape drives, increasing the network bandwidth and employing D2D2T backup approaches or alternatively backing up to disk in a remote location using a high bandwidth connection. In any case, the backup window is a serious problem as soon as the archive exceeds a few Terabytes.

The use of RAID replication for data protection is becoming increasingly popular as the cost of magnetic disk capacity decreases. This approach avoids the need for a backup window but has two main limitations when it comes to data protection – it does not provide protection in case of accidental data deletion and it requires that the replica RAID is in a different location. The primary and replica data should be located such that fire and other major disasters will not destroy both data sets.

Limitations of Optical Disk Systems

Optical disk systems have been popular for medical image archiving but suffer from problems of poor reliability, low data transfer rates, relatively low capacity per disk and high cost.

Poor reliability has been associated with CD and DVD media formats because the optical disks are not protected within a cartridge, making them prone to damage and data loss. This is overcome by using the more expensive formats such as MO, UDO and ProData Disk (PDD) which employ a protective cartridge.

All optical disk systems suffer from low data transfer rates. Even the latest formats have transfer rates that are very low relative to those that are easily achieved with magnetic disk and tape. For example, it is difficult to achieve more than 2 MB/s when writing to the inner diameter of a UDO disk. Whereas, modern tape formats and hard disks routinely will write at 20 MB/s or higher.

The highest capacity optical disk formats available today are UDO and ProData Disk. UDO offers 15 GB per side on a double sided disk and PDD provides 23 GB on a single sided disk. These are relatively low compared to the latest magnetic disks (over 300 GB) and tape formats such as LTO-3 (400 GB) and SAIT (500 GB). Consequently, an optical disk based system requires many more units of media compared to a disk-tape based system. This is not really much of a problem for low capacity storage requirements below a terabyte. However, for multiple terabytes, it is problematic. For example, a data store based on PDD disks requires over 400 units of media for 5 TB including a duplicate copy for disaster recovery. This in turn increases the difficulty of media management and makes for bulky and expensive optical library systems.

¹ Please contact XenData for details of technology partners that can provide DICOM connectivity.

High Capacity Tape for Data Archiving

Today’s leading tape formats provide reliable long term archiving and offer native capacities from 100 GB to 500 GB per cartridge:

Format	Capacity	Cartridge Type
AIT-3	100 GB	8 mm
AIT-4	200 GB	8 mm
SDLT-600	300 GB	Half inch
LTO-3	400 GB	Half inch
SAIT	500 GB	Half inch

A half inch format tape cartridge is illustrated opposite. Both 8 mm and half inch cartridges provide strong physical protection for the media.



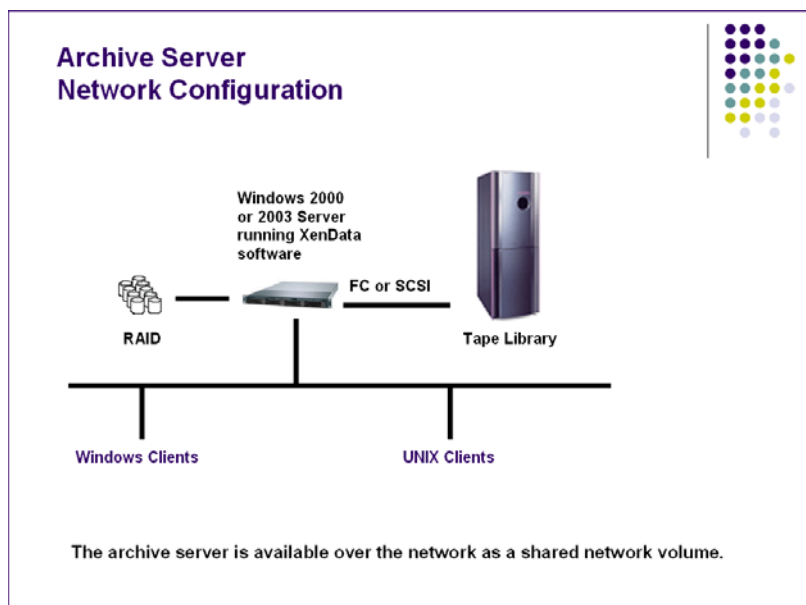
Projected archival lifetime varies from manufacturer to manufacturer and is typically specified with lifetimes up to 30 years. For example AIT and SAIT tapes from Sony use high stability Advanced Metal Evaporation techniques in the manufacture of their media which provides a 30 years projected lifetime.

Medical Image Archive Server Configuration

XenData Archive Series software virtualizes a magnetic disk logical drive (typically on RAID) and a tape library as a single server logical drive which may be shared over the network. A typical network configuration is shown in the diagram opposite. Windows clients access the archive as a shared logical drive letter. UNIX clients may also access the server via NFS.

The server operating system may be either Windows 2000 Server or Server 2003. In addition the equivalent Windows NAS operating systems are supported.

The magnetic disk logical drive may be direct attached RAID or provided from a SAN. For smaller installations, it could even be a partition on an attached magnetic disk.



The system administrator defines policies that determine where data files are physically stored. These policies support hierarchical storage management and automatic tape cartridge replication. They define file retention times on RAID, the specific group of tapes to which the file will be written, scheduling of replica tape cartridge updates, etc. A single server may have many different policies, tailored to the needs of the different file types that are being archived. In all cases, when a file is written to tape, it is first written to magnetic disk and then immediately written to tape once the file is closed.

The file management policies are defined by the system administrator via the XenData Management Console which is a Microsoft Management Console (MMC) snap-in. This provides a familiar administration interface which can be easily integrated with other MMC snap-ins.

The software supports both rewritable and Write-Once-Read-Many (WORM) tape cartridges. WORM tape provides an unalterable record of the data written to it and provides best practice when authenticity of the digital content is important. XenData Software supports the simultaneous use of WORM and rewritable media within the library.

XenData Software Key Functionality

XenData Archive Series Software provides the following key functionality:

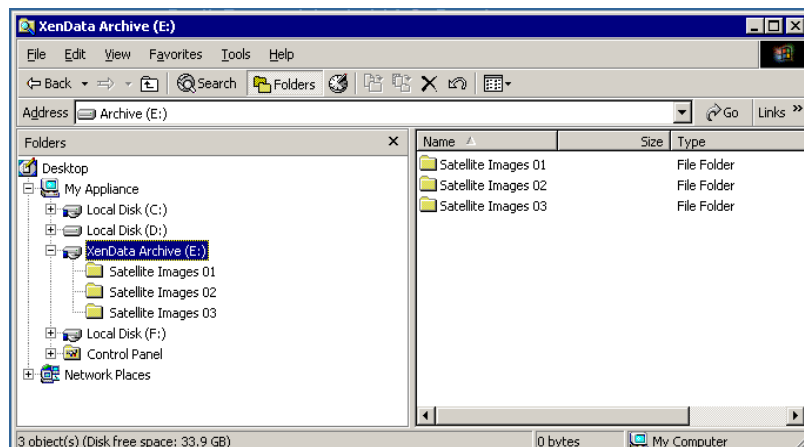
- **Standard File System.** An archive server managed by XenData software works with standard applications because the RAID and tape library are presented as a standard drive letter that is shared over the network.
- **HSM based on De-Migration.** Administrator-defined policies determine where files are physically stored. The XenData approach uses hierarchical storage management (HSM) based on file de-migration and it delivers both data protection and high performance, taking full advantage of the attributes of both magnetic disks and tape.
- **Continuous Backup.** Files are continuously written to disk and tape, immediately generating two physical instances of the data. This approach avoids the major limitations of conventional backup including the need for a backup window.
- **Automatic Update of Tape Cartridge Replicas.** Additional data protection is provided by the automatic generation of tape cartridge replicas.
- **WORM Tape Support.** The software provides a full audit trail of the file system. When WORM tape is used, the system meets the most stringent legal requirements for digital document retention.

Standard File System

XenData software presents RAID and tape as a single rewritable logical drive letter. Even when using WORM tape cartridges, files written to the XenData logical drive may be 'deleted' or modified as though they are being written to a standard rewritable magnetic disk. However, all versions and all deleted files are retained on tape and can be accessed via a XenData utility.

An example of how the XenData logical drive letter appears to Windows Explorer is shown in the illustration opposite. In this case, drive letter E is controlled by the XenData software.

The use of a standard rewritable file system provides a major benefit; it means that the XenData archive works with standard applications without modification.




Hierarchical Storage Management based on De-migration

XenData software provides policy-based HSM based on file de-migration. The use of file de-migration, rather than classic file migration, means that multiple instances of a file are generated which in turn provides the data protection benefits described in the next section.

Classic HSM migrates infrequently used files from magnetic disk to a lower cost medium, leaving a pointer on the magnetic disk. This conventional approach is less relevant today because of the low cost of disk storage. In addition, it creates problems because the data on magnetic disk must still be backed up.

HSM based on de-migration takes a different approach which allows integration of data protection, replacing conventional backup. With XenData software, the HSM and data protection are driven by administrator defined policies. A typical policy is illustrated in the diagram below.



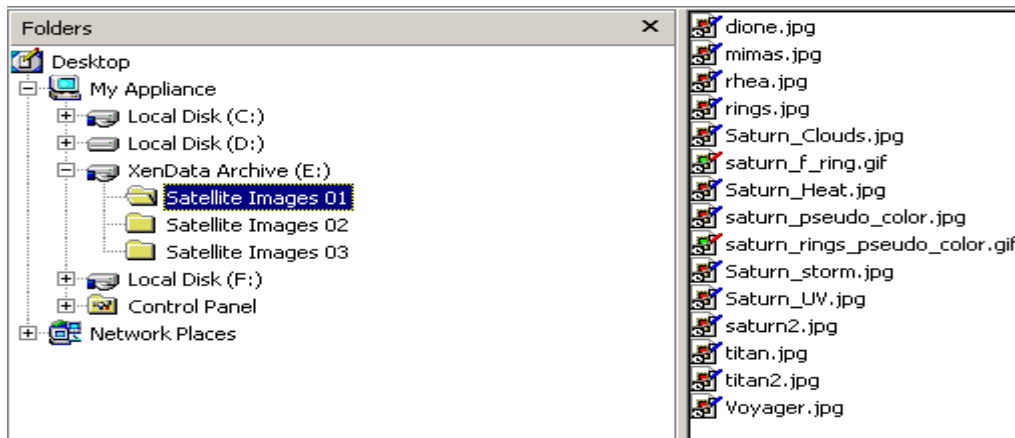
With the file management policy illustrated above, a file would be written to both RAID and a specified group of tape cartridges. This immediately provides two copies of the file for data protection purposes. The system administrator will have predefined a replication schedule and the number of tape copies for replication. After all copies of the file have been produced, the file may be de-migrated from RAID according to the administrator-defined policy. On reading the file, it is always retrieved from the fastest physical storage device. If it is stored on magnetic disk and tape, the file will be read from disk. If the file is available only on tape, it will be retrieved from tape.

The main levels of file storage hierarchy supported by XenData Software are:

- **On-Line** with one instance on RAID and at least one instance on tape. In this case the file will be retrieved from RAID when accessed over the network.
- **Near-Line** with at least one instance on tape within the library. Typically the file will be written to more than one tape cartridge using XenData's automatic tape cartridge replication capabilities.
- **Off-Line** with one or more instances on tape, all of which have been exported from the tape library.

De-migrating (i.e. removing) a file from RAID will cause it to move from on-line to near-line, freeing up space on the magnetic disk. The rules for de-migration are defined by the administrator and may be based on the time elapsed since a file was written or last read. For example a file may be de-migrated 7 days after it was written so long as it has not been read in the last 24 hours. If it is then read from tape, it will be restored to RAID and retained on-line until 24 hours have elapsed without it being read.

When a file is de-migrated from RAID and is on either near-line or off-line tape, an attribute is set which indicates that the file is no longer on-line. This changes network timeout periods to allow retrieval of the file from media with long access times. It also changes the appearance of a file within Windows Explorer – a small clock is added to the bottom left of the file icon as shown below.



Data Protection via Continuous Backup and Tape Replication

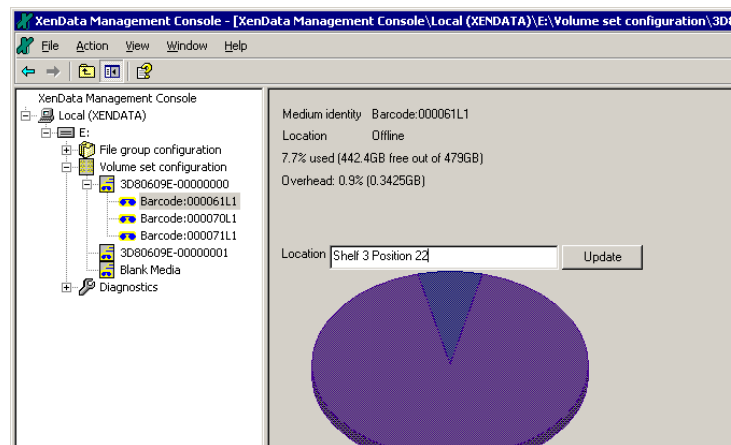
The XenData storage solution replaces conventional backup because it automatically produces multiple instances of files and metadata, as defined by administrator policies. The multiple instances of a file are created in two steps:

1. At the time of writing to the XenData archive, the file is written to both magnetic disk and tape.
2. Then the file is written to additional tape cartridges according to tape replication policies defined by the administrator.

Step 1 provides a continuous backup of files which overcomes the two major limitations of conventional backup:

- There is no backup window - as soon as a file has been written to magnetic disk and the file is closed, it is immediately written to tape.
- On recovery, all data is restored unlike conventional backup where changes to the file system since the last backup are lost.

The automatic replication of tape cartridges - step 2 described above - is implemented in a way that supports rotation of tape cartridges between an off-site location and the tape library. There is no need to wait until a tape cartridge becomes full before it can be replicated. When a tape cartridge replica is reintroduced into the library after being retained off-site, it is automatically brought up to date by the system.



The XenData Management Console allows the Administrator to identify the location of off-line tape cartridges.

Data recovery is very easy to accomplish whether it is restoration of the complete file system or the recovery of an individual file that was accidentally deleted. Also, it is easy to test that files have really been written to tape, unlike conventional backup. An individual file can be quickly retrieved from a tape cartridge and opened using the XenData History Explorer utility which provides confidence that the system is functioning correctly.

<input checked="" type="checkbox"/>	Throughput XTN001.pdf	88049	Tue Jul 22 2003 18:15:01	4
	Throughput.doc	46592	Fri Mar 21 2003 09:04:19	4
	Throughput.xls	24064	Fri Mar 21 2003 09:04:40	4
	XenData Archive Series Data Sheet.pdf	279303	Wed Feb 12 2003 18:09:13	4

e		Status		
2003 11:04:00	Not on RAID	Archived (Barcode:000024)	fetch from tape	
2003 11:04:00	Not on RAID	Archived (Barcode:000024)	fetch from	
2003 11:04:00	Not on RAID	Archived (Barcode:000024)	fetch from	
2003 11:04:00	Fetch from RAID	Archived (Barcode:000024)	fetch from	

XenData History Explorer lists all instances of all versions of all files, including deleted and renamed files. Any instance of a file may be opened or saved to another location.

WORM Tape Support

XenData software can simultaneously manage both Write-Once-Read-Many and rewritable tape cartridges in the same tape library. Policies that are set by the Administrator define which files are written to which type of tape cartridge.

WORM is both non-erasable and non-rewritable. Consequently, WORM tape cannot be reformatted and after data is written to the tape, it cannot be changed. However, further data can be appended.

XenData software manages this non-erasable medium in such a way that files can be 'deleted' within the file system in the normal way. However, they are really being hidden from the normal file system interface and can be retrieved using the XenData History Explorer utility. The History Explorer utility can be used to identify and restore all old file versions and all deleted files. When using WORM media, this capability to identify and retrieve all deleted files and old versions is vital when legal compliance is an important issue.

The combination of XenData software and WORM tape ensures data authenticity. It provides legal compliance and best business practice, including meeting HIPAA data authentication requirements.

Additional Features

Additional features of an archive server based on XenData software include:

- Multiple Tape Sets - the Administrator can group related files together on the same set of tapes.
- Dynamic Expansion of Tape Sets - the system will dynamically expand tape sets to meet capacity demands.
- Cartridge Spanning - the system supports storage of large files across multiple tape cartridges.
- Security - XenData software is fully integrated with the Microsoft Windows security model, based on Active Directory.
- Standard Tape Format - open standard POSIX TAR file format is used on the tape, allowing the tape cartridges to be read using standard third party utilities.
- Familiar administration of policies via Microsoft Management Console.
- Comprehensive diagnostics tightly integrated with the Windows operating system.

In Summary...

XenData Archive Series Software creates a high capacity storage solution on a Windows server which combines hierarchical storage management, integrated data protection and file version management, making it ideal for multi-TB medical image archives.

The XenData solution eliminates the major problem associated with using RAID and conventional backup, in that it does not require any backup window. Relative to an optical disk system, it offers higher reliability than bare optical media, higher performance, lower total cost of ownership and much higher scalability.

For more information, please contact:

XenData: ■ *USA +1 925.472.6522* ■ *UK +44 1223 370114* ■ *Germany +49 89 99216 422*