

Implementing a Digital Video Archive Based on the Sony PetaSite and XenData Software



The Video Edition of XenData Archive Series software manages a Sony PetaSite tape library on a Windows Server 2003 platform to create a digital video archive that is ideal for the demanding requirements of the broadcast industry.

This Sony - XenData solution is well proven with successful installations in North America and Europe. It offers high performance, writing and reading at many times real time. Yet it is non-proprietary, presenting the digital archive as a standard Windows file system which allows it to be used by multiple applications. In addition, the video archive is highly scalable from terabytes to over a petabyte per server.

This paper discusses system architecture, the S-AIT digital tape format, PetaSite configuration and server hardware, video file management options, as well as maintenance and training considerations.

1. Introduction

Broadcasting is rapidly moving from video tape to a file-based environment. It is generally accepted that digital automation and delivery offer better performance and lower costs than traditional analogue approaches. When it comes to video archiving, retaining the historical material as files and storing those files on magnetic disk arrays (RAID) and digital tape also offers the same benefits.

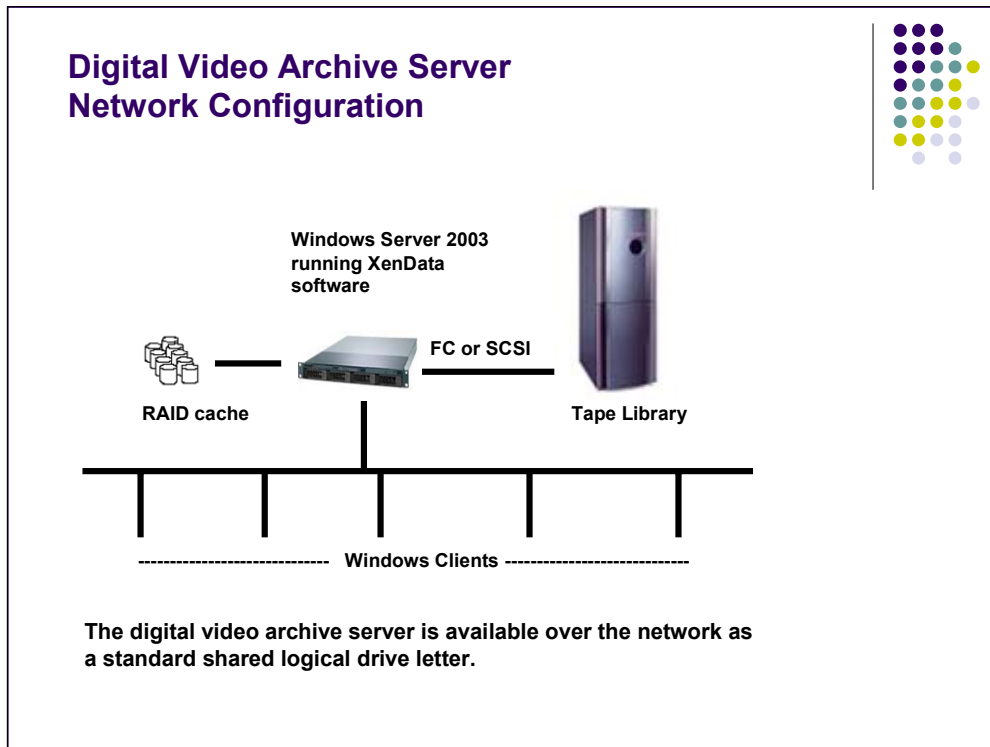
The Video Edition of XenData Archive Series software manages a Sony PetaSite and RAID cache on a Windows Server 2003 platform, creating a high performance digital video archive. The XenData solution will typically write and retrieve video files at over 200 Mbits per second, representing data transfers that are much faster than real time.

The solution is very well integrated into the Windows server environment, presenting the PetaSite as a standard Windows file system. This non-proprietary approach means that the archive can be used by multiple applications and it does not tie the user to any particular asset management or automation solution.

2. System Architecture Considerations

Digital Video Archive Network Configuration

A digital video archive server running XenData software presents the PetaSite tape library and RAID cache as a single Windows logical drive letter. The combined storage within the tape library and RAID effectively appear as a very large capacity magnetic disk. This means that no special interface is required and the video archive will work with multiple standard applications without modification.



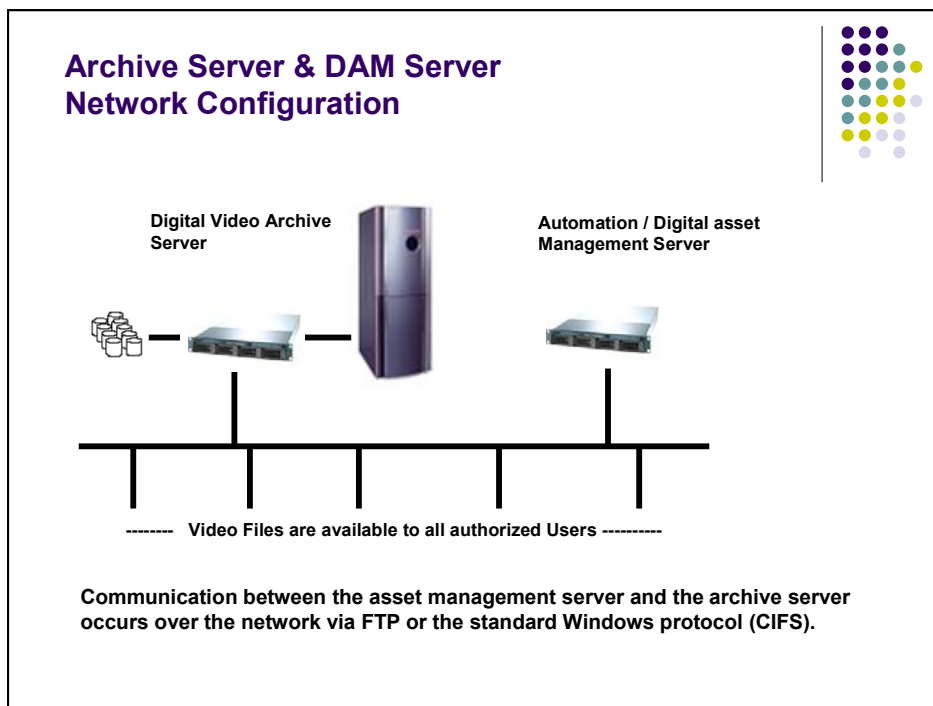
A basic network configuration is shown above. The digital tape library is connected to the server via SCSI and / or Fibre Channel interfaces. The RAID cache may be implemented any number of ways, for example as direct attached SATA disk arrays or from a SAN.

The Video Edition of XenData software is optimised for transferring large files over the network via FTP or the standard Windows network protocol, CIFS. Other network protocols are supported, but FTP and CIFS are recommended for high performance digital video applications.

Combining Asset Management & the Digital Archive

Digital Asset Management (DAM) provides indexing of digital material and the ability to search and retrieve the assets of interest. The asset management system stores the indexed material as files which are held in one or more digital video archives.

XenData software creates a digital video archive, managing a PetaSite tape library and RAID cache and presenting these physical storage devices as a standard Windows file system. Furthermore the software provides hierarchical storage management, data protection via tape cartridge replication, partial read capabilities and file security.



The DAM and digital video archive may each have a dedicated server, as illustrated in the configuration shown above. Alternatively, the DAM and video archive may reside on the same server.

3. Hardware Options

S-AIT Digital Tape Format

Until recently, specialist tape formats were required to meet the performance requirements of professional digital video. This is no longer the case, as the S-AIT format which was developed for the general IT industry, more than meet these performance needs. S-AIT digital tape cartridges use very stable materials and offer a media archival lifetime of 30 years. This represents a huge advance compared to conventional video tapes which have much shorter lifetimes.

S-AIT (Super Advanced Intelligent Tape) was developed by Sony and has a strong market share in digital video applications. At 500 GB, it offers the highest capacity per cartridge. Using a recording bit rate of 25 Mbps, this is equivalent to 44.4 hours per cartridge.

S-AIT characteristics are summarized below.

S-AIT Tape Format	S-AIT-1
Capacity per cartridge (Native capacity i.e. without compression)	500 GB
Equivalent hours recording per cartridge at 25 Mbps	44.4 hours
Equivalent hours recording per cartridge at 50 Mbps	22.2 hours
Maximum Data Transfer Rate in Megabytes per second (without compression)	30 MB/s
Media Archival Lifetime	30 years

More information about this format is available from www.sait.com.



The 500 GB capacity of an S-AIT tape cartridge is equivalent to over 44 hours recorded at 25 Mbps.

Tape Library Options

XenData software supports the S-PetaSite tape library range which starts at 60 tape slots (equivalent to over 2,600 hours at 25 Mbps) and can be expanded to approximately 3000 tape slots (over 130,000 hours at 25 Mbps). In addition for smaller requirements, XenData software supports the Sony CSM-20 which provides a capacity of 20 slots and includes one or two S-AIT drives.

The number of tape drives within the library is an important consideration and a major factor in the tape library cost. Considerations are discussed below:

- For small archives with only one tape drive within the library such as might be configured with the Sony CSM-20, it is important to be very careful in setting file management policies to

prevent the single drive becoming a bottleneck degrading performance. When only one drive is present, we strongly recommend that a large RAID cache is used; only one set of tapes is configured and that tape replication is scheduled to occur at a time when there is no other writing or reading activity.

- With two tape drives, the concern over tape drive bottleneck is much less severe than for a single drive system. Even so, we recommend that only one tape set should be configured.
- Libraries with three or more drives are recommended for larger archives as are typical with the S-PetaSite. XenData software will intelligently manage many drives, allowing simultaneous writing and multiple file accesses. With three or more drives, the administrator can configure multiple tape sets and can group related files together on the same set of tapes, if required.

Server Hardware

With large files written to a server running XenData software, the data transfer rates for both writing and reading are typically limited only by network speed or by the tape drives. For example, with S-AIT drives, 1 GB Ethernet and a 2.4 GHz dual processor, sustainable transfer rates of over 200 Mbits/s are easily achievable.

The minimum server hardware requirements are available from www.xendata.com. Even a modest specification server will deliver high performance. However, it is important that the server has sufficient RAM. A minimum of 1 GB is required, but 4 GB is recommended.

The required capacity of the RAID cache is very dependent on the HSM and replication policies that are set by the system administrator. However, in most implementations at least 1 TB of cache is recommended.

4. Choosing File Management Policies

HSM and Data Protection Administrator Settings

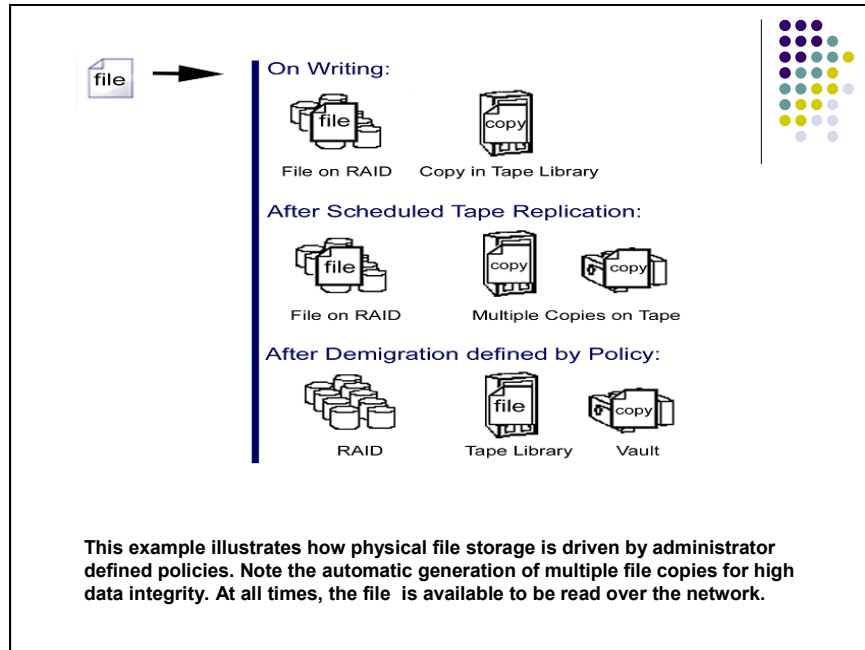
The system administrator defines policies that determine where data files are physically stored on the digital video archive. These policies support hierarchical storage management (HSM) and automatic tape cartridge replication.

XenData Archive Series software supports three main levels of storage hierarchy:

- Online with one instance of a file on RAID and, in addition, there will typically be one or more instances on tape. In this case the file will be retrieved from RAID when accessed over the network.
- Near-line with at least one instance of a file on tape within the library and no instance on RAID. When a near-line file is accessed over the network, the XenData software automatically transfers the file from tape to RAID cache. As soon as the file transfer to RAID starts, the file is also transferred over the network.
- Off-line with no instance on RAID and one or more instances of a file on tape, all of which have been exported from the tape library.

Data protection is achieved by automatically generating multiple instances of a file. The XenData software can automatically produce copies of digital tapes for off-site retention.

A single server may have many different policies, tailored to the needs of the different file types that are being archived. A typical XenData file management policy is illustrated in the diagram below.



Keeping Proxies Online

XenData software stores the video files on RAID and digital tape according to policies set by the administrator. Policies determine how long files will be retained on RAID and these RAID retention rules may be tailored to the different file types that are being archived. This means that low resolution proxies can be kept permanently online on RAID, whereas the high resolution broadcast media can be held near-line on digital tapes within the library.

Enabling Partial Read

XenData software manages very large files by using controlled file fragmentation. The administrator can optionally define policies that split large files into multiple fragments. This is performed by the XenData software in a way that is hidden from the applications that are reading and writing files and it is particularly useful for multi-gigabyte files. With digital video applications there is often a need to read only a portion of a very large file. For example, consider reading a portion a 20 GB file that has been archived with a policy that splits it into forty 500 MB fragments. In this case, when a portion of a file is being read from tape, the software instructs the drive to rapidly seek to the start of the first fragment that contains the required portion of the file. The system then only retrieves the fragments that contain the requested data. Without the controlled fragmentation provided by the XenData software, the complete 20 GB file would have to be read from tape, which would take many minutes. In practice, this approach of using controlled file fragmentation is very easy to implement and greatly enhances performance when dealing with large files.

Setting File Permissions

XenData Archive Series software integrates fully with the Microsoft Windows security model, based on Active Directory. Files and directories have administrator-definable security attributes just as they do with standard Microsoft file systems and access control checks are performed in the same way. This means that no special or proprietary set up is required for configuring file permissions.

5. Operational Considerations

Tape Library Maintenance

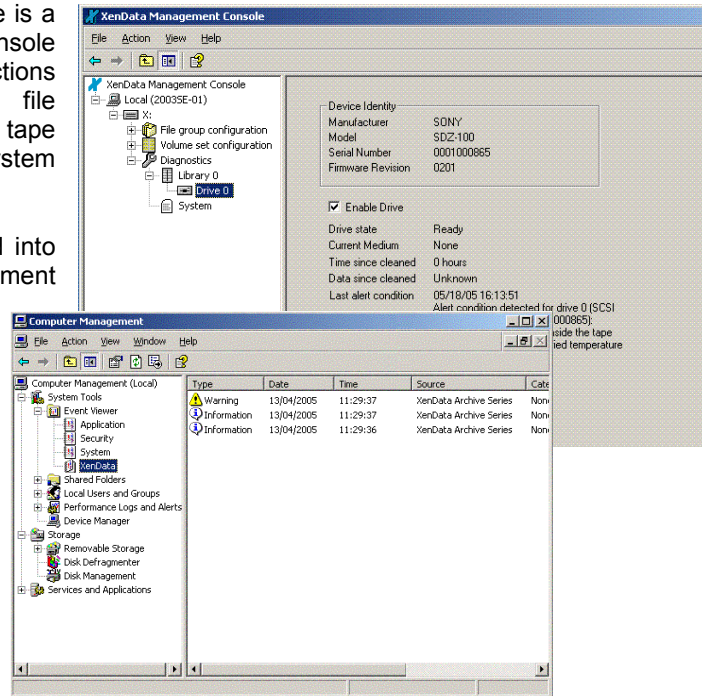
XenData Archive Series software monitors alerts provided by the drives within the tape library and logs warning and error conditions in the Microsoft event log and XenData Management Console as appropriate.

Periodically, drives require cleaning using a tape cleaning cartridge. We recommend that at least one tape cleaning cartridge is held within the tape library. The cleaning cartridge will be located within a regular tape cartridge slot and when cleaning is required, the drive will issue an alert which prompts the XenData software to automatically undertake a cleaning cycle.

System Training

XenData Archive Series software is very tightly integrated into the Microsoft server operating system:

- The main administration console is a Microsoft Management Console snap-in. It has three main sections addressing the setting of file management policies, tape management policies and system diagnostics.
- Event logging is fully integrated into the Microsoft computer management system. Event logs appear in a dedicated XenData Event Viewer section.
- File and directory security attributes are defined just as they are with standard Microsoft file systems.



Consequently for an administrator familiar with Windows Server 2003, the training required is only a few hours.

Glossary

CIFS 'Common Internet File System'. It is the standard protocol used by Windows computers to communicate on a network.

DAM 'Digital Asset Management'.

Fibre Channel A technology for transmitting data between computer devices at data rates of up to 2 Gbps. Fibre Channel is especially suited for connecting computer servers to shared storage devices and for interconnecting storage controllers and drives.

FTP 'File Transfer Protocol'. FTP is a protocol commonly used to copy files between two computers on the Internet. Both computers must support their respective FTP roles - one must be an FTP client and the other an FTP server.

HSM 'Hierarchical Storage Management'. XenData Archive Series software supports three main levels of storage hierarchy: online, near-line and off-line.

NAS 'Network Attached Storage'. XenData Archive Series software may be installed on NAS file servers running Windows Storage Server 2003.

RAID 'Redundant Array of Independent Disks'

RAM 'Random Access Memory'

S-AIT 'Super Advanced Intelligent Tape'. This is a popular mid-range tape format developed by Sony.

SAN 'Storage Area Network'. The magnetic disk logical drive managed by XenData software may be provided from a SAN. Note that it must be formatted with NTFS and configured as a dynamic disk.

SCSI 'Small Computer System Interface'. This is a common interface for connecting computers and peripherals.

Tape Alert Tape drives and other hardware employ a mechanism called Tape Alert to send maintenance information to external devices. XenData Archive Series software periodically monitors Tape Alert information and informs the user or takes other appropriate action where necessary.

TAR A term derived from 'Tape ARchive'. It is a popular open standard data format optimized for archiving to tape and is the format used by XenData Archive Series software.

For more information, please visit www.xendata.com or contact XenData:

■ USA +1 925.472.6522

■ UK +44 1223 370114

■ Germany +49 89 99216 422