


Implementing an Automated Digital Video Archive Based on the Video Edition of XenData Software



The Video Edition of XenData Archive Series software manages one or more automated data tape libraries on a Windows Server 2003 platform to create a digital video archive that is ideal for the demanding requirements of the broadcast industry.

The solution is 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, choice of digital tape format, tape library and server hardware, video file management options, as well as maintenance and training considerations.

1. Introduction

Digital video archiving in broadcast has conventionally involved complicated architectures, proprietary formats and proprietary interfaces. This has resulted in solutions which have been expensive to install and difficult to maintain.

In contrast, the Video Edition of XenData Archive Series software creates a digital archive on a Windows Server 2003 platform with a straightforward architecture and non-proprietary interfaces. The software is tightly integrated into the Windows server operating system which means it delivers high performance with a simple and elegant configuration.

The XenData archive has a standard file system interface appearing as a single Windows logical drive letter. The solution is optimized for use with the standard Windows network protocol (CIFS) or FTP file transfers. This non-proprietary approach to the interface 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.

In addition to its standard file system interface, the solution offers many other non-proprietary features:

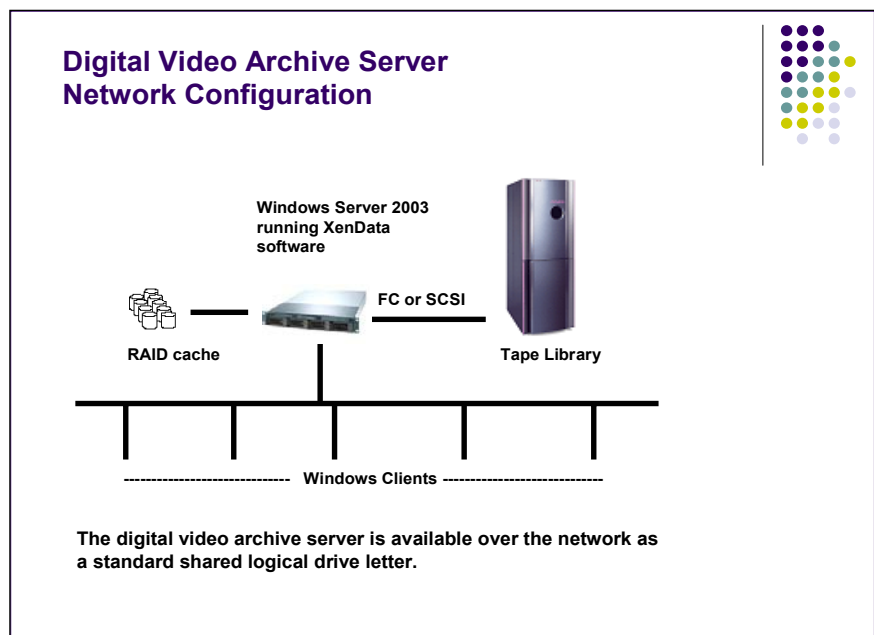
- All file types can be archived on the system and partial file restore has been implemented in a way that is not specific to the file type.
- The system uses the open standard POSIX TAR format for recording to data tape.
- All informational, warning and error messages are logged in the standard Windows Event Log and may be sent as on-screen messages or e-mail alerts using the XenData Alert Module.
- The system fully complies with the Microsoft security model based on Active Directory, which means that tedious special administration of file permissions is not necessary.

2. System Architecture Considerations

Digital Video Archive Network Configuration

A digital video archive server running XenData software presents the digital 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.

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



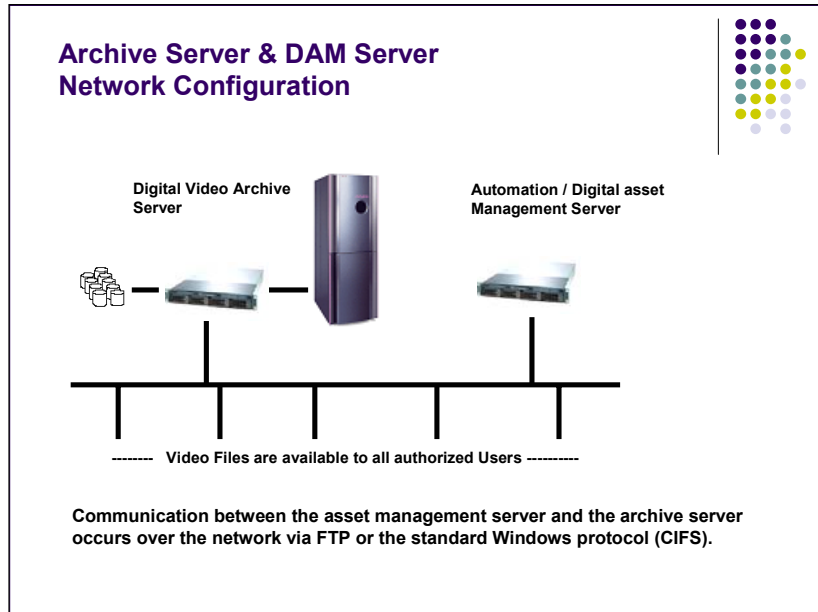
The Video Edition of XenData software is optimised for transferring large files over the network via FTP or the standard Windows network protocol, CIFS. Other network protocols are supported, but FTP and CIFS are recommended for high performance digital video applications.

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 digital 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, for smaller system, the DAM and video archive may reside on the same server.



3. Hardware Options

Choosing a Data 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 data tape formats developed for the general IT industry more than meet these performance needs. All of the leading tape formats use very stable materials and are available with a media archival lifetime of 30 years. This represents a huge advance compared to conventional video tapes which have much shorter lifetimes.

There are a number of alternative data tape formats which are well suited technically for digital video. However, the dominant market share is held by the LTO (Linear Tape Open) format, which is also known as Ultrium. LTO tape drives are manufactured by multiple vendors and the most popular for use in automated tape libraries are manufactured by HP and IBM. More information is available from www.lto-technology.com.

At the time of writing this paper, the LTO format is transitioning from its third to fourth generation - from LTO-3 to LTO-4. The new LTO-4 drives will both read and write LTO-3 tape cartridges.

LTO tape cartridge characteristics are described below.

	LTO-3	LTO-4
Capacity per cartridge (Native capacity i.e. without compression)	400 GB	800 GB
Equivalent hours recording per cartridge at 25 Mbps	35.5 hours	71 hours
Equivalent hours recording per cartridge at 50 Mbps	17.7 hours	35.5 hours
Maximum Data Transfer Rate in Megabytes per second (without compression)	80 MB/s	120 MB/s
Media Archival Lifetime ¹	30 years	30 years



The 800 GB capacity of an LTO-4 tape cartridge is equivalent to over 71 hours recorded at 25 Mbps.

Tape Library Options

LTO tape formats are available in tape libraries from a large number of suppliers including ADIC, Hewlett Packard, IBM, Overland Storage, Qualstar, Quantum, Rorke Data, Sony and Spectra Logic. A current list of the tape libraries supported by XenData is available from www.xendata.com.

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

- For small archives with only one tape drive within the library, it is important to be very careful in setting file management policies to prevent the single drive becoming a bottleneck and in turn 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 active tape set should be configured.
- Libraries with three or more drives are recommended for large archives. 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.

¹ 30 years is typical for LTO media. However, for specific brands, please check with your vendor.

Server Hardware

With video 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 magnetic disk performance. For example, with LTO-3 or LTO-4 drives, 1 GB Ethernet and a 2.4 GHz dual processor with 4 GB of RAM, sustainable transfer rates of over 300 Mbits/s are easily achievable.

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.

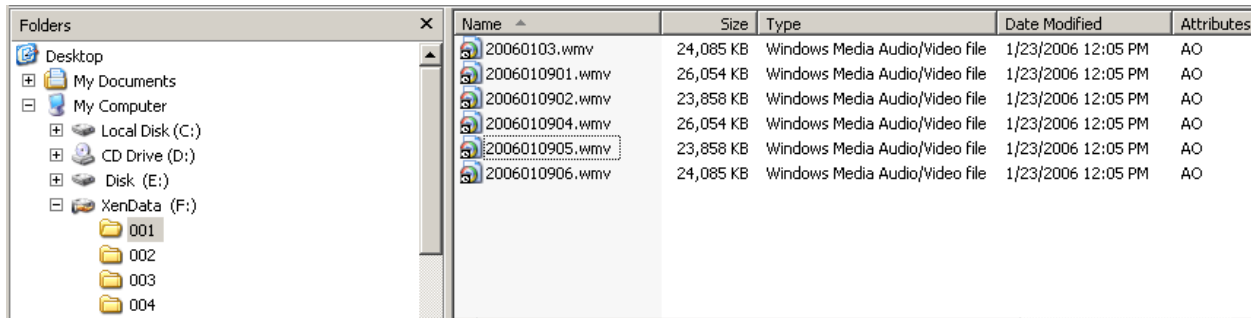
Data protection is achieved by automatically generating multiple instances of a file. The XenData software can automatically produce copies of digital tapes for off-site retention.

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 opposite. On writing a file, it is first written to RAID. As soon as the file has been successfully written to disk, it is put into a queue to be written to a primary tape cartridge. After completion of this operation, there are two instances of the file – one on disk and one on tape.

Tape cartridge replication is optional and is scheduled according to an administrator policy. For a library with sufficient tape drives, it may be scheduled to occur immediately or within a defined time period. Alternatively, it may be scheduled to occur daily at a specific time. Following replication, the file is written on one or more duplicate tape cartridges and it becomes eligible for deletion (also termed 'flushing') from RAID.



After deletion from RAID, the offline attribute bit is set and files are still available from tapes within the library. The Microsoft offline bit 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.



On reading from tape, a file is automatically restored to RAID as it is simultaneously transferred over the network.

Keeping Proxies Online

XenData software stores the video files on RAID and data 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 of a 40 GB file that has been archived with a policy that splits it into forty 1 GB 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 40 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

Periodically, drives require cleaning using a tape cleaning cartridge. We recommend that at least one 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.

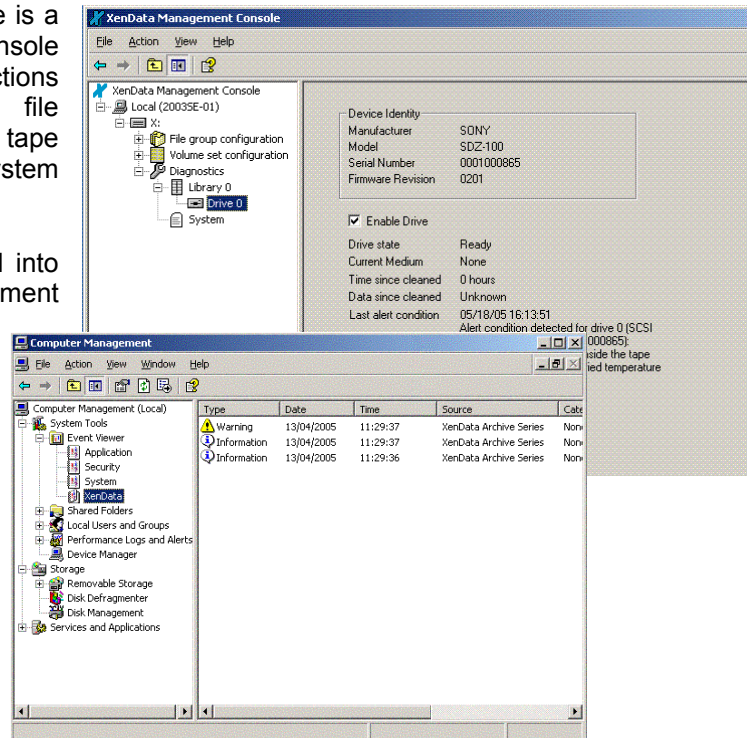
XenData Archive Series software monitors alerts provided by the tape library and tape drives and provides a comprehensive set of informational, warning and error conditions in the Microsoft Event Log. These are very useful for obtaining early warning of any hardware or media problems.

The Microsoft Event Log messages are converted into e-mail and on-screen alerts using the XenData Alert Module which is companion product. The Alert Module has five pre-configured categories of events that address the needs of an archive operator and system support personnel. For example, a message indicating that the tape library is running out of blank media can be sent to the operator. Whereas, a message indicating faulty operation of a tape drive can be addressed to appropriate IT support personnel.

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 4 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.

LTO 'Linear Tape Open'. This is the most popular high performance data tape format which is also known as Ultrium.

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'

SAIT 'Super Advanced Intelligent Tape'. This is a data 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.

Ultrium High performance tape format, also known as LTO.

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