

# **XenData6 Server Software Administrator Guide**

**Version 6.11**

## Contents

### 1. Introduction

- 1.1. Overview
- 1.2. Hierarchical Storage Management
- 1.3. Windows Offline Attribute
- 1.4. LTO Cartridge Formats: LTFS and TAR
- 1.5. Alternate Data Stream Support
- 1.6. Barcode Management
- 1.7. Installations including a Tape Library
- 1.8. Installations with only Stand-Alone Tape Drives
- 1.9. Disk Requirements
- 1.10. Compatible Operating Systems
- 1.11. Antivirus Compatibility
- 1.12. License Administration

### 2. Concepts

- 2.1. Introduction
- 2.2. File Groups
- 2.3. Volumes, Volume Sets and Automatic Replication
- 2.4. Partial File Restore and Cartridge Spanning
- 2.5. Pending Write Mode
- 2.6. Offline File Management
- 2.7. File Version Management
- 2.8. File Security
- 2.9. Handling of File Delete and Rename Operations
- 2.10. Tape Catalogs and Finalization
- 2.11. Recovering Tape Space using Repack

### **3. Administering the System**

#### **3.1. Overview of Implementing File Management Policies**

- 3.1.1. To Start the XenData Management Console
- 3.1.2. To Navigate Around the XenData Management Console
- 3.1.3. To Display XenData Management Consoles for other Archives on the Network
- 3.1.4. New Installation Check List

#### **3.2. Configuring Volume Sets**

- 3.2.1. Understanding Volume Sets
- 3.2.2. About the Blank Media Set
- 3.2.3. About the Quarantined Media Set
- 3.2.4. About Volume Set Defaults for a New Installation
- 3.2.5. To Configure a Volume Set
- 3.2.6. To Define Replication for a Volume Set
- 3.2.7. To Define Replication Timing
- 3.2.8. To Allocate Cartridges to a Volume Set
- 3.2.9. To Create a New Volume Set
- 3.2.10. To Replace a Missing Replica Cartridge
- 3.2.11. To Delete a Volume Set

#### **3.3. Managing Tapes**

- 3.3.1. About Rewritable Cartridges
- 3.3.2. About WORM Cartridges
- 3.3.3. About Importing Cartridges
- 3.3.4. About Reformatting Cartridges
- 3.3.5. To Reformat a Rewritable Cartridge
- 3.3.6. To Write Files to Tape
- 3.3.7. To Read Files Directly from Tape
- 3.3.8. To Display File Statistics for a Volume
- 3.3.9. To Display Information About a Cartridge
- 3.3.10. About Write-Protecting Cartridges
- 3.3.11. To Write-Protect a Cartridge
- 3.3.12. About Building and Rebuilding Tape Catalogs
- 3.3.13. To Build a Tape Catalog
- 3.3.14. To Rebuild a Tape Catalog
- 3.3.15. To Repack a Volume
- 3.3.16. To Cancel a Repack Operation
- 3.3.17. About Re-tensioning Tapes
- 3.3.18. To Re-tension a Tape
- 3.3.19. About Verifying the Data on a Tape
- 3.3.20. To Verify Data on a Tape
- 3.3.21. About Finalizing a Tape
- 3.3.22. To Finalize a Tape
- 3.3.23. To Export a Cartridge
- 3.3.24. To Set the Administrator Defined Information for a Cartridge

#### **3.4. Allocating Files to File Groups**

- 3.4.1. About File Groups
- 3.4.2. To Allocate Files to a File Group
- 3.4.3. To Construct the "File Name or Path Pattern" or "Exclude Pattern"
- 3.4.4. To Create a New File Group
- 3.4.5. To Rename a File Group
- 3.4.6. About Changing the Order of File Groups
- 3.4.7. To Change the Position of a File Group in the List

### 3.4.8.Examples of Allocating Files to File Groups

## 3.5. Setting Policies for File Groups

- 3.5.1.About Storage Options for a File Group
- 3.5.2.To Select Storage Options for a File Group
- 3.5.3.File Fragmentation
- 3.5.4.Disk Retention Rules
- 3.5.5.File Group Advanced Options

## 4. Using Mac Clients

- 4.1. About Using Mac Clients
- 4.2. File Group Policies
- 4.3. Disabling Icon Preview
- 4.4. Disabling Alternate Data Streams
- 4.5. Support of OS X Characters

## 5. File Explorer and Windows Explorer Extensions

- 5.1. About File Explorer and Windows Explorer Extensions
- 5.2. Smart Copy and Paste
- 5.3. Flushing
- 5.4. Pre-fetching
- 5.5. Volume View
- 5.6. History Explorer
- 5.7. Enhanced Properties via XenData Tab

## 6. Metadata Backup

- 6.1. Overview
- 6.2. Using the Program
  - 6.2.1.Starting the Program
  - 6.2.2.Selecting Backup or Restore
  - 6.2.3.Making a Pre-defined Backup
  - 6.2.4.Making a Custom Backup
  - 6.2.5.Restore from backup

## 7. Scheduler

- 7.1. Overview
- 7.2. Using the Program
  - 7.2.1.Starting the Program
  - 7.2.2.To Add a Task
  - 7.2.3.To Edit a Task

- 7.2.4.To Delete a Task
- 7.2.5.To Immediately Run a Task
- 7.2.6.To Immediately Stop a Task

### **7.3. About the Scheduler Display**

### **7.4. How to Configure a Metadata Backup Task**

### **7.5. How to Configure a Deferred Write Task**

## **8. System Recovery**

### **8.1. System Rebuild from Tape**

- 8.1.1.About File System Rebuild from Tape
- 8.1.2.To Rebuild a File System from Tape

### **8.2. In Case of Tape Drive or Library Failure**

- 8.2.1.To Disable a Tape drive or Library
- 8.2.2.To Re-enable a Tape Drive or Library
- 8.2.3.Options In Case of Tape Library or Stand-Alone Drive Failure

## **9. Diagnostics & Maintenance**

### **9.1. XenData Event Log**

### **9.2. XenData System Trace Log**

- 9.2.1.To Open the Trace Log
- 9.2.2.To Change the Level of Detail in the Trace Log
- 9.2.3. Automatic Generation of Trace Files

### **9.3. To View Tape Drive or Library Diagnostic Information**

### **9.4. Tape Drive Cleaning**

- 9.4.1.About Tape Drive Cleaning
- 9.4.2.To Manually Clean a Drive

### **9.5. Tape Cartridge Alert State**

- 9.5.1.About the Tape Cartridge Alert State
- 9.5.2.To Clear the Alert State for a Tape

## **10. Report Generator**

### **10.1 About the Report Generator**

### **10.2 Starting the Report Generator**

### **10.3 Creating, Saving and Restoring Reports**

- 10.3.1 To Create a Report
- 10.3.2 To Save a Report
- 10.3.3 To Display a Saved Report

## **10.4 File Search Reports**

- 10.4.1 To Run a File Search Report
- 10.4.2 Interpret a File Search Report

## **10.5 Recoverable Space Reports**

- 10.5.1 To Run a Recoverable Space Report
- 10.5.2 Interpret a Recoverable Space Report

## **10.6 Tape Contents Reports**

- 10.6.1 To Run a Tape Contents Report
- 10.6.2 Interpret a Tape Contents Report

## **10.7 UnArchived Files Reports**

- 10.7.1 To Run an UnArchived Files Report
- 10.7.2 Interpret an UnArchived Files Report

## **10.8 Volume Contents Reports**

- 10.8.1 To Run a Volume Contents Report
- 10.8.2 Interpret a Volume Contents Report

## **11. Glossary**

## Chapter 1. Introduction

### 1.1 Overview

XenData6 Server is the 6th generation of XenData Archive Series software. It runs on a Windows computer and provides archiving to tape cartridges tailored for the needs of high performance digital video applications. The software uses a dedicated magnetic or solid state disk volume and manages data tape drives and tape libraries, creating a highly scalable digital archive.

XenData6 Server software provides the following key functionality:

- **Standard File System Interface** The archive appears locally as a Windows file system within a single logical drive letter. Files are stored in a file and folder structure. The archive file system represents the files stored on the dedicated magnetic or solid state disk volume, on tape cartridges in the attached tape libraries and tape drives and on offline tape cartridges. The archive logical drive letter or any folder within it may be configured as a network share.
- **Hierarchical Storage Management** Administrator-defined policies determine where files are physically stored. Three tiers of HSM are supported: file instances may be stored on the dedicated disk volume, near-line tape and offline tape.
- **Disk Retention Determined by Policy** The disk retention time is determined by administrator-defined policies. The standard Windows offline file attribute identifies when a file is only available from data tape and no longer on the disk volume.
- **Standard Network Protocols** The solution supports CIFS/SMB and FTP file transfers.
- **Windows and Mac Client Compatibility** Windows and Apple OS X clients are natively supported without need for loading any client software.
- **Choice of LTFS or TAR Formats** The archived files are written to LTO tape cartridges using either the open standard TAR format or the LTFS exchange format. LTFS is applicable to LTO-5 and later generations of rewritable tape cartridges.
- **Support for Multiple Cartridge Groups** The software allows groups of files to be allocated to specified groups of tape cartridges. The administrator-defined policies can be used to group related files together on the same set of cartridges.
- **Dynamic Expansion of Cartridge Groups** The system will dynamically expand cartridge groups to meet capacity demands.
- **Automated Cartridge Replication** The software automatically generates replica data tape cartridges that may be exported from the library for off-site retention. Furthermore, the cartridges may be rapidly imported into a replica DR system.
- **Intelligent Barcode Management** The software selects tape cartridges in barcode order and also automatically selects matched barcodes for replicated cartridges.
- **Repack of Cartridges** This operation copies current files to new cartridges, excluding deleted files and old file versions. After completion, the original tape cartridges are quarantined and may be reformatted as blank. Repack is useful for recovery of cartridge capacity due to deleted files and old file versions. It is also useful for migrating files to later generations of cartridges.
- **Partial Read of Large Files** With very large files there is often a need to read only a portion of the file. For example, this frequently occurs with large video files when a short clip is requested. XenData6 Server software supports partial reading of large files based on byte offset. For a total solution that provides partial restores based on time code, a third party application that supports partial restores based on time code should also be used.

- **Cartridge Spanning** The administrator-defined policies can be set to allow or prevent individual files being spanned across multiple tape cartridges. This option is particularly useful when very large files are being archived.
- **Optimized Restores** The system restores a queue of files in the shortest possible time. The restore requests are processed in an order that minimizes unnecessary tape movement.
- **Self-Describing Data Tapes** Each tape cartridge contains all the file system metadata necessary to recover all the files stored on it whether using TAR or LTFS.
- **Transfer of Content between Systems** Export and import functions allow cartridges to be easily transferred from one location to another.
- **Microsoft Security** The archive file system is fully integrated with the Microsoft Windows security model based on Active Directory.
- **WORM Tape Support** In addition to standard rewritable tape cartridges, the software supports unalterable WORM data tapes when using the TAR tape format.
- **File Version Control** The software provides comprehensive file version control. Deleted files and old file versions may be restored from tape (unless the files have been purged using a repack operation).

This version of XenData6 Server software automatically installs three companion XenData utilities:

- **Metadata Backup** The XenData metadata backup program backs up and restores the archive file system metadata and the XenData Management Console settings. It allows rapid restore of the system in case of failure of the magnetic disk storage and is described in Chapter 6.
- **Scheduler** The XenData Scheduler can be used to schedule metadata backups and to defer writing to tape cartridges for scheduled periods. It is described in Chapter 7.
- **Report Generator** The report generator creates archive system reports including a report showing the contents of any tape, including offline tapes and is described in Chapter 10.


## 1.2 Hierarchical Storage Management

XenData6 Server software supports three levels of hierarchical storage:

- **Online** with one instance of a file on the disk under XenData control and, in addition, there may be one or more instances on tape. In this case the file will be restored from disk.
- **Near-line** with at least one instance of a file on tape within a tape drive or library and no instance on disk. Often the file will be written to more than one tape cartridge using XenData's automatic tape cartridge replication capabilities.
- **Offline** with no instance on disk and instances of a file on one or more tape cartridges, all of which have been exported from the available tape drives and libraries.



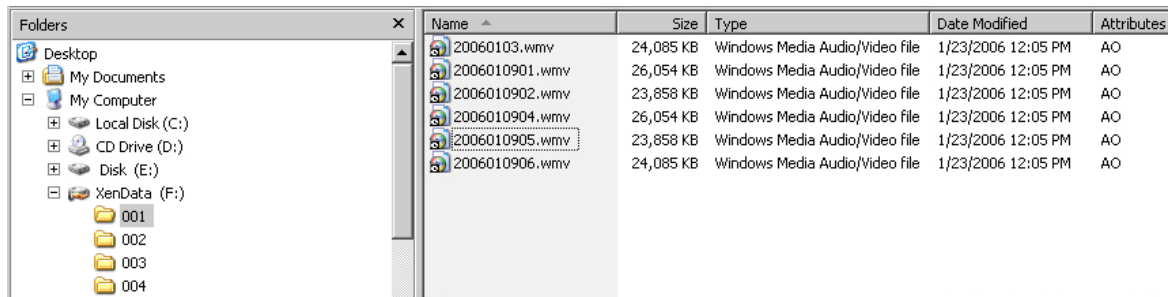
The administrator establishes disk retention rules that determine when a file instance is to be removed from disk. The action of removing a file instance from disk is termed ‘flushing’. Once the disk retention rules have been established, the software manages everything automatically. A typical storage policy is illustrated below.



**Note:** The magnetic or solid state disk volume under XenData control is often referred to in XenData documentation as the disk cache.

### 1.3 Windows Offline Attribute

XenData6 Server software uses the standard Windows offline file attribute to identify files that are no longer online on disk storage. After a file has been successfully written to tape and is no longer on disk, the offline attribute is set. This Microsoft file attribute changes network timeout periods to allow retrieval of files from storage media with long access times. It also changes the appearance of a file within Windows Explorer for Windows versions prior to Windows 8 and Windows server 2012 - a small clock or cross is overlaid on the file icon as illustrated below.



The state of the offline file attribute for the three tiers of HSM is as follows.

<b>HSM Tier</b>	<b>Description</b>	<b>Offline Attribute</b>
<b>Online</b>	One instance of the file is on disk and, in addition, there may be one or more instances on tape.	Not set
<b>Near-line</b>	At least one instance of the file is on tape within a tape library or tape drive and no instance is on disk.	Set
<b>Offline</b>	There is no instance of the file on disk and there are one or more instances on tape, all of which have been exported from the tape libraries and ejected from stand-alone drives.	Set

#### 1.4 LTO Cartridge Formats: LTFS and TAR

This version of XenData6 Server supports both LTFS (Linear Tape File System) and TAR (Tape ARchive) cartridge file system formats. These formats define how data is written to the tape media: LTFS and TAR use different data structures for the file data and file system metadata that are written to tape. When configuring a Volume Set (a pool of tape cartridges), the administrator selects either TAR or LTFS as the cartridge file system format.

The LTFS format was developed by IBM and announced in 2010. Since then, it has been widely adopted, making it an exchange standard which allows cartridges to be moved between systems created by different vendors. LTFS uses two partitions on the LTO cartridge: a small index partition for maintaining the contents of the tape and a large data partition for the stored files. It is applicable to rewritable LTO-5 and later generations of LTO cartridges. A tape cartridge written using the LTFS format is self-describing and the contents of the cartridge can be determined quickly by reading an index partition on the tape.

The following LTO drive and cartridge combinations support LTFS:

- LTO-6 drives and 2.5 TB LTO-6 rewritable cartridges
- LTO-6 drives and 1.5 TB LTO-5 rewritable cartridges
- LTO-5 drives that support dual partitions and 1.5 TB LTO-5 rewritable cartridges

The TAR format was introduced in 1979 and is a widely adopted open standard supported by many operating systems including most versions of UNIX, Linux and Microsoft Windows Services for UNIX. It is applicable to all data tape types including all generations of rewritable and WORM LTO. A tape cartridge written using the TAR format is self-describing. However, when using a basic TAR implementation, the whole tape must be scanned to determine its contents. XenData6 Server software extends the TAR format by maintaining a catalog of the cartridge contents and writing this catalog to the end of the tape in a finalization operation. When a finalized tape cartridge written using the TAR format is moved to another XenData system, the contents are quickly determined because the system automatically reads the catalog.

## 1.5 Alternate Data Streams Support

Alternate data streams, also known as 'NTFS streams' and 'named streams', are additional data streams that can be associated with a file name. This version of XenData6 server software supports alternate data streams when writing to a Volume Set configured with the TAR format.

Alternate data streams are supported by Windows NTFS volumes, Mac OS X Server from version 10.5 and some NAS (Network Attached Storage) operating systems. They are used in a number of ways including the following.

- Mac OS X clients from version 10.6 automatically enable alternate data streams over SMB when connected to a Windows NTFS share or the share of a volume managed by XenData6 Server. The stream contains file metadata. The streams may be disabled on the Mac client, in which case the client writes 'AppleDouble' files consisting of a data fork file and resource fork file.
- Windows Internet Explorer will add a stream named 'Zone-Identifier' to files downloaded from the Internet. Windows uses this data for security purposes.

Alternate streams are not listed in Windows Explorer, and their size is not included in the file's size. Streams may be listed on a Windows computer from the command prompt using the 'DIR' command with a '/R' parameter. The following filename format is used: "filename:streamname" (e.g., "text.txt:extrastream").

## 1.6 Barcode Management

Barcode labels are available for all data tape cartridge formats supported by XenData6 Server software and are strongly recommended for keeping track of tapes when using a tape library. The barcodes are readable by both humans and a tape library that includes a barcode reader. (Data tape libraries typically include a barcode reader as standard.)

Most data tapes, including LTO, also include an in-cartridge memory chip. When a tape cartridge is initialized in a tape library under XenData control, the barcode information is written to the in-cartridge memory chip.



This is particularly useful when the tape has been exported from the library and inserted into a stand-alone tape drive. Tape drives do not contain barcode readers but can read the contents of the in-cartridge memory. XenData6 Server software provides a consistent identification of tape barcode for all tapes that were originally initialized in a tape library, even when a tape is being used in a stand-alone tape drive.

In addition to automatic recording of the barcode to the in-cartridge memory chip, XenData6 Server software matches barcodes for replicated tapes and selects tapes in barcode order.

The inventory of blank media provided by the tape library is sorted by the XenData software in alphanumeric barcode order. The software identifies a matched set as barcodes which differ by only one letter (e.g. "A" and "B") in any one of the barcode characters. The lowest matched set is always allocated for replicated sets of tapes. For non-replicated tape sets or when no matched set exists, the tapes are simply allocated in alphanumeric order.

For example, if we had the following sequence of tape barcodes in the blank media set:

**X0003AL4**  
**X0005AL4**  
**X0007AL4**  
**X0007BL4**  
**X0008AL4**  
**X0008BL4**  
**X0009AL4**  
**X0009BL4**

The system would next allocate X0007AL4 and X0007BL4 to a duplicated set of tapes. However, if the system were allocating a tape to a non-replicated set, X0003AL4 would be selected.

## 1.7 Installations including a Tape Library

XenData6 Server software may be configured and licensed to manage one or more tape libraries, optionally combined with one or more stand-alone tape drives.


A Windows file server with a tape library is illustrated in the diagram opposite. The tape library is typically connected via Fibre Channel (FC) or Serial Attached SCSI (SAS) to the archive server. A fibre channel library may be connected directly to FC ports on the archive server or via a fibre channel switch.

All archived files appear within one Windows logical drive on the server whether those files are on the disk cache, on a tape within the library or are on an offline tape. This single logical drive may be shared over the network, making the archive accessible to network clients, via CIFS/SMB or FTP.

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

- For archives with only one tape drive within the library, it is important to be careful in setting the file management policies to prevent the single drive from becoming a bottleneck and in turn degrading performance. When only one drive is present, we recommend that frequently accessed files are retained online on the disk cache and that tape replication is scheduled to occur at a time when there is no other writing or reading activity.
- A system with two tape drives within the library is a good choice for many archive applications when there is not a high volume of restore operations.
- Libraries with three or more tape drives are recommended for large or frequently accessed archives. XenData software will intelligently manage many drives, allowing simultaneous writing and multiple file accesses.

All installations that include a tape library, even if it has only one drive within it, support tape cartridge replication.




## 1.8 Installations with only Stand-Alone Drives

XenData Archive Series software may be configured and licensed to manage one or more stand-alone tape drives on a Windows server.

A Windows file server with one data tape drive is illustrated in the diagram opposite. In a configuration with only one tape drive, setting the Pending Write Mode option (as described in Chapter 2) is often desirable because it allows the disk cache to act as a buffer when writing a set of files to multiple tapes.

Even if there were two tape drives attached to the archive server, all archived files appear within one Windows logical drive on the server whether those files are on a tape within one of the drives or are offline.



## 1.9 Disk Requirements

The Windows server must be configured with at least two logical drives on magnetic or solid state disk - a boot drive (usually C:) and another drive dedicated for use by the XenData software. The XenData software is installed on the boot drive but it requires minimal available capacity. The dedicated disk logical drive is used to store file system metadata, for read and write caching of files held on tape and to store all files that are retained online. This dedicated logical drive must be formatted with NTFS and must be configured as a dynamic disk.

Capacity requirements for the dedicated logical drive vary considerably depending on how the system is configured. The total capacity requirement is the sum of the following.

- **File System Metadata** Each file in the archive file system requires two NTFS clusters. If the dedicated logical drive is formatted using the default cluster size of 4096 bytes, a file system with one million files will require 8 GB for the metadata.
- **Read and Write caching** The dedicated logical drive provides caching for each file that is being written and all files that are open due to being read by a program. Using a capacity allowance of twenty times the largest file size is adequate for most environments.
- **Online Files** Any group of files may be retained online. The Administrator uses the XenData Management Console to define policies for online retention. If tape cartridge replication is not employed, the system can be configured to flush files immediately after writing or reading, in which case the additional capacity requirements for online files is zero. At the other extreme, the system can be configured to keep all current files online, in which case the additional capacity requirements for online files is the sum of all the file sizes within the file system.

**Example Calculation** for an installation with only stand-alone tape drives is given below:

- Up to a 250,000 files - this requires 2 GB for metadata
- Maximum file size of 5 GB - 100 GB is allowed for caching
- Immediate file flushing and no replication - zero allowance is required for online retention

A total dedicated logical drive capacity of 102 GB or higher is estimated for this installation.

**Example Calculation** for a 200 TB digital video archive using a tape library is given below:

- 100,000 files with high resolution content and 100,000 low resolution proxy files - this requires 1.6 GB for metadata
- Average high resolution file size of 2 GB - 40 GB is allowed for caching
- Average low resolution proxy file is 100 MB and all must be retained online - this requires 10 TB for online retention.

A total dedicated logical drive capacity of just over 10 TB is estimated for this installation.

## 1.10 Compatible Operating Systems

This version of XenData6 Server software is certified for installation on Windows Server 2012 R2, Windows Server 2008 R2 and Windows 7 X64 Professional Edition.

## 1.11 Antivirus Software Compatibility

When installing antivirus protection on an archive running XenData6 Server, it is important to choose an antivirus (AV) solution that has been certified. XenData6 Server software and AV software use file system filtering techniques and there may be undesirable interactions if you use an AV product that has not been certified.

Please refer to XenData Tech Note XTN1201, available from the support section of the XenData web site ([www.xendata.com](http://www.xendata.com)), for information about certified AV products.

## 1.12 License Administration


XenData6 Server software is usually licensed for a particular tape library and drive configuration. A XenData Activation Code is required to run the software and this enables the chosen tape hardware. Licensing is administered via the XenData License Administration Program.

### To start the XenData License Administration Program

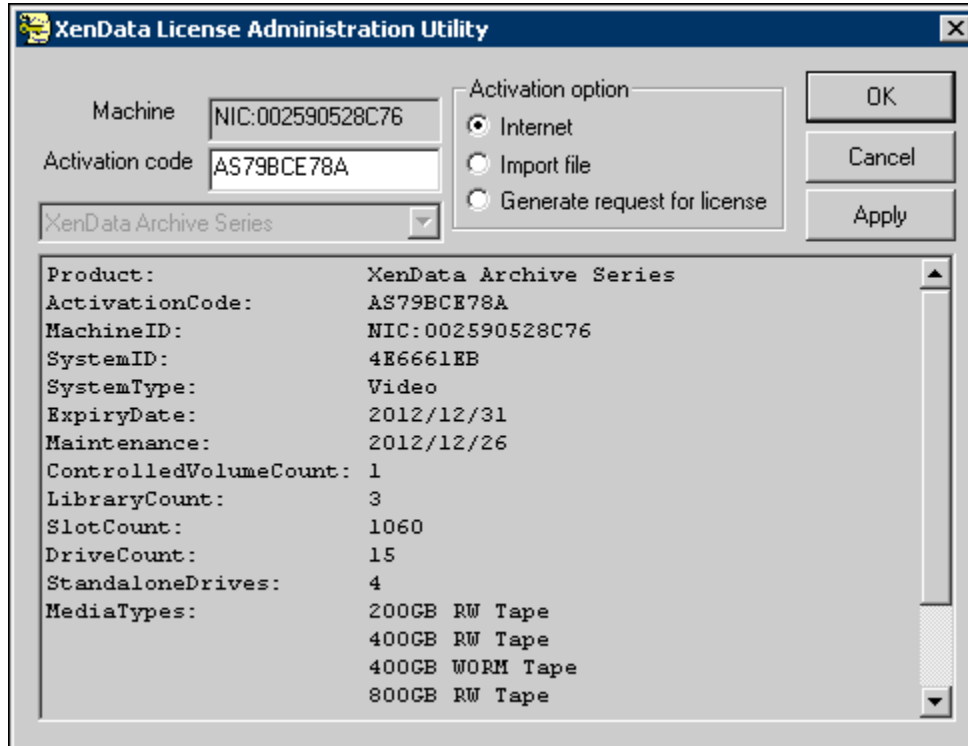
For Windows Server 2008 R2 and Windows 7:

1. Click **Start**
2. Click **Programs**
3. Click **XenData**
4. Click **License Administration**

For Windows Server 2012 R2:

1. Click on the **Start** button, located in the lower-left hand corner of the screen
2. Click the arrow  near the lower-left corner of the Start screen
3. Click on **License Administration** shown under the XenData category

The license administration program detects the unique network interface card (NIC) ID of the system and prompts for a XenData Activation Code. A digitally signed activation file must then be obtained from a XenData license server. There are two automated ways to obtain this via the Internet as described in the first two options below. If there is no Internet access at all, the digitally signed activation file can be obtained from XenData technical support as described in option 3 below.



### Option 1

When Internet connectivity is available to the target server, select **Internet** as the Activation Option and then select **Apply**. A digitally signed activation file is generated for the system and transmitted over the Internet back to the target server. This process usually takes a few seconds and completes the licensing process.

### Option 2

When Internet connectivity is not available to the target server but is available on another machine, select **Generate request for license** as the Activation Option and then select **Apply**. A "Save As" window is then displayed. Choose a file name and location for the file and select **Save**. Copy this file to the machine that has Internet access; open the file and follow the instructions. A license key file is generated. Then copy the saved license key file back to the target server, open the license administration utility and select **Import file** as the Activation Option and then select **Apply**. Browse to the saved license key file and select **Open**. This completes the licensing process.

### Option 3

If Internet access is not available, send us an email at [support@xendata.com](mailto:support@xendata.com) with the XenData Activation Code and the file created using the **Generate request for license** as described in Option 2 above and a digitally signed activation file will be returned by email. Save the file to a local drive, select **Import file** as the Activation Option and click **OK**. The License Administration program prompts for the license file: browse to the correct file and select **Apply**. This completes the licensing process.

## Chapter 2. Concepts

### 2.1 Introduction

XenData6 Server software enables the system administrator to define policies that allow:

- groups of files (termed 'File Groups') to be defined based on the file name and folder;
- groups of files to be stored on defined sets of tape cartridges (termed 'Volume Sets');
- groups of files to be retained online on disk for predefined periods;
- temporary retention of files on disk in case of no writable tape cartridges being available;
- tape format, either TAR or LTFS;
- automatic generation of tape cartridge replicas;
- fast partial restore of very large files;
- spanning of files across tape cartridges;
- notification of requests for files that are stored offline;
- file version management; and
- configuration of file security to control file writing and access.

This functionality is based on concepts that are defined and discussed in this chapter.

In addition, the following topics are discussed:

- the method used to keep track of the contents of individual tapes using tape catalogs; and
- the repack operation which recovers tape space from deleted files and old versions of files.

### 2.2 File Groups

A File Group is a collection of files that have the same file management policy and consequently are all treated in the same way by the system. Every file that is handled by the system belongs to exactly one File Group. Files are assigned to a File Group on the basis of name and path. This assignment can be based on the name of the folder that contains a file, the name of the file or both. For example, the system could be configured with the following three File Groups.

- File Group 1 contains all files with names of the form \*.tmp, regardless of folder and these are saved to disk but are not saved to tape.
- File Group 2 contains all other files saved to the folder \project1\ and these are permanently retained online on disk and are also saved to tape to a specific Volume Set called 'Tapes01'
- File Group 3 contains all other files saved to the folder \project2\ and these are retained online on disk for 24 hours after they were written or last read and they are also saved to tape to another Volume Set called 'Tapes02'

A file or folder cannot be written to the system if a File Group is not present that matches its file or folder name and path. In the example above, files can only be stored if written to the folders \project1\ or \project2\ or if they have a name of the form \*.tmp. Attempts to write other files will not be permitted. If required, the system can be configured with an additional 'catch all' File Group that sets a policy for all other files.

File Group characteristics are defined by the administrator using the XenData Management Console. The following parameters must be defined for each File Group:

- File name or path pattern for the File Group. This, together with the relative position of the File Group in the XenData Management Console display, defines the files that are allocated to this File Group.



- The position of the File Group relative to other File Groups in the XenData Management Console display, which defines the order in which file name or path pattern selection is applied. If a file name matches the selection criteria for more than one File Group, it will be assigned to the first matching File Group.
- Enable or disable 'Save files to tape'.
- Selection of a Volume Set, if the File Group is saved to tape.
- Enable or disable file fragmentation. If enabled, the fragment size must be set. File fragmentation must be enabled to allow partial file restore and cartridge spanning.
- Retention periods on magnetic disk, if the File Group is saved to tape.

The policies defined by the administrator determine how files are stored on disk and tape. The policies do not change folders or files within the Windows file system. For example, folders cannot be created by using File Group rules; new folders are created by using standard utilities like Windows Explorer.

### 2.3 Volumes, Volume Sets and Automatic Replication

XenData6 Server software can automatically create multiple tape cartridge replicas. The XenData term 'Volume' refers to a complete set of replica tapes which, when up-to-date, all contain the same data. If replication is not enabled, a Volume refers to an individual tape cartridge.

A Volume Set stores files from designated File Groups and consists of one or more Volumes. As more data is written to a Volume Set, the initial Volume will eventually become full. At a preset threshold, defined by the administrator, the system will automatically add another Volume by taking the appropriate number of tapes from the blank media set and will extend the Volume Set.

The data on replicated tapes in a Volume Set are kept synchronized whenever the tapes are available to the system. If one or more tapes in a Volume Set are removed from the library, the system maintains a record of which files need to be written to those tapes. When tapes are reintroduced into the library, the data on them is automatically brought up to date.

With the exception of the blank media set and quarantined media set, all media within a Volume Set must be either WORM tape or rewritable. Tape replicas must all be of the same capacity. One unique Volume Set, termed the blank media set, contains all the tapes that are recognized by the system but are not formatted for storing data. These may be new (unused) tapes or rewritable tapes that have been reformatted using the XenData Management Console.

Another unique Volume Set, termed the quarantined media set, contains all cartridges that have been imported into the library but for some reason are currently unusable by the system. Typically, this will be because:

- a cartridge has previously been repacked; or
- a cartridge has been used by a different application (such as a backup application) and inserted into a stand-alone tape drive or imported into the tape library; or
- because an error occurred while the system was trying to identify the contents of the cartridge.

Following configuration of the File Groups, Volume Sets and any associated replication requirements, the system operates completely automatically. Files written to the logical drive under XenData control are automatically allocated to File Groups. Files allocated to File Groups with 'Save files to tape' enabled have an assigned Volume Set and are automatically written to both disk and tape. If replication is enabled for the assigned Volume Set, this occurs automatically.

### 2.4 Pending Write Mode

When all tapes in a Volume Set become full or unavailable, the system response to an attempt to write additional files to the Volume Set depends on the Volume Set configuration settings. All tapes in a Volume Set may become full or unavailable for a number of reasons:

- use of stand-alone tape drives and the inserted tape cartridge or cartridges have simply become full;
- use of a tape library with insufficient blank tape cartridges;
- a failure of the tape drive or library hardware; or
- a tape cartridge media error.

If the 'Write to disk if no writable media are available' option has been enabled and all tapes in a Volume Set become full or unavailable, the system automatically enters the Pending Write Mode and will accept more data. However, if the 'Write to disk if no writable media are available' option has not been enabled, the system will not accept any more data and will report "disk full" when an attempt is made to write to the Volume Set.

When the system enters the 'Pending Write Mode', it defers writing to tape and continues writing to disk. When free space within the Volume Set becomes available again, the system automatically 'catches up' and writes the deferred files to tape. When the system is in the Pending Write Mode, a comprehensive set of warning messages are sent to the Windows Event Log. These include notification of entering and leaving the Pending Write Mode and running short of space in the disk cache. When the 'Write to disk if no writable media are available' option is enabled, we recommend that the XenData Alert Module be configured to provide notification via email and / or on-screen message of these warning messages.

## 2.5 Partial File Restore and Cartridge Spanning

The ability to quickly restore a portion of a file from tape is often useful in professional video when a short video clip is being read from a very large file. It might take many minutes to restore the whole file from tape and consequently the ability to partially restore can greatly improve the performance of the whole system.

The ability to partially restore a file is obtained by enabling file fragmentation for a File Group. Furthermore, file fragmentation allows a large file to be spanned across multiple Volumes. File fragmentation causes the XenData software to 'chunk' files as they are written to the archive. When a fragmented file is restored, it is automatically concatenated.

Enabling file fragmentation typically provides benefits for File Groups that predominantly consist of multi-gigabyte files. If file fragmentation is enabled, the system has the following characteristics:

- when a portion of a file is read from tape, only the applicable fragments will be read, saving both transfer time and space on the disk cache;
- when an application modifies a large file by appending, the appended data will be written to tape as one or more additional fragments, saving space on the tape cartridge;
- if an application modifies a small part of a large file, for example by updating an index at the beginning of the file, then only the fragments containing modified data will be written to tape; and
- on writing a file, the system may span multiple Volumes if the File Group advanced option to permit file spanning is enabled and, when spanning occurs, complete fragments of spanned files will be written to each of the spanned Volumes.

If file fragmentation is not enabled, the system has the following characteristics:

- when a file or portion of a file is read from tape, the complete file will be read from tape;
- when a file is modified, the new version of the file will be completely written to tape; and
- on writing a file, the system will always write the complete file to a single Volume.

## 2.6 Offline File Management

XenData6 Server software can be configured to provide three tiers of storage hierarchy:

- Online with one instance of the file on disk and typically one or more instances on tape. In this case the file will be retrieved from disk when accessed.
- Near-line with no instance on disk and at least one instance on tape within a tape drive or library. Near-line files are retrieved automatically when accessed.
- Offline with no instances on disk and one or more instances on tape, all of which have been ejected from the tape drives and exported from the tape libraries. To access an offline file, it must be brought near-line by manually importing a tape cartridge that contains the file into a tape library or alternatively inserting it into a tape drive.

Offline files appear in the Windows file system but when they are accessed by a program, a message is returned that identifies that the file is not available. Also, the XenData6 Server software puts a message into the XenData area of the Windows Event Log that identifies which tape cartridges contain the file. When the XenData Alert Module is installed, on-screen messages and e-mail alerts are also generated that identify the file name and the tapes that contain the file. The XenData Alert Module is a companion product that is licensed separately.

## 2.7 File Version Management

The Windows file system interface provides access to the latest version of a file and does not permit access to old file versions or to deleted files. However, if old file versions or deleted files were saved to tape, they can be accessed using History Explorer, which is extended functionality within Windows Explorer that is provided by XenData6 Server software.

If a file is updated with a newer version by overwriting or appending, the XenData software assigns a new version number. A file's version number increases by one every time it has data written to it. Note that the version number does not increase for every individual write operation, just for every file open that is followed by a write. Version 0 of a file never contains any data; the first time an application writes to the file, the version number is incremented to 1.

If a file is deleted and then a new file of the same name containing data is created, the system starts again with version 1 of the new file and a new generation is created. For example, if a new file is created for the first time on a system, it will have generation 0 and version 1. If that file is then deleted and another file of the same name is saved, this new file will be designated as generation 1 and version 1.

## 2.8 File Security

XenData6 Server software integrates fully with the Microsoft Windows security model, based on Active Directory. Files and directories have user-definable security attributes just as they do with standard Microsoft file systems and access control checks are performed in the same way. The security model is extended to deleted files and old versions of files made available to users via History Explorer, or another application written using the XenData6 Server API. In these cases, the security allocated to prior versions of a file or directory is the same as that applied to the most recent version, regardless of the security applied when the old version was originally in use. This feature allows system administrators to update access controls for old files based on changing business requirements.

## 2.9 Handling of File Delete and Rename Operations

To make each archival cartridge (e.g. an LTO cartridge) self-contained, even when a file is deleted or renamed after the cartridge becomes nominally full, a change in the way delete and rename operations are handled was made from version 6.11 of XenData6 Server software.

In versions prior to 6.11, the delete or rename record is written to the current writable volume which is not necessarily the volume to which the file was written. From version 6.11 of XenData6 Server these records are written to the volume that contains the last version of the file. The applicable cartridge(s) must be available for writing i.e. mounted in a standalone drive or library under XenData control. This applies to both TAR and LTFS formatted LTO cartridges. If the cartridge on which the file was last written is offline, it will be impossible to rename or delete the file and a message will be written to the Windows Event Log stating that the required cartridge(s) is offline and giving the cartridge barcode(s). After the cartridge is made available near-line, you will be able to perform the rename or delete operation.

If the file is written to a replicated volume, then only one of the replica cartridges needs to be near-line for the delete or rename to be successful. In this case, the offline replica cartridge will be identified in the XenData Management Console as 'Needs updating', and will be updated when it is put back into the library or inserted in a stand-alone drive.

If a file that is fragmented and spans more than one cartridge is deleted or renamed, the delete or rename record will be written to all cartridges that contain that file.

## 2.10 Tape Catalogs and Finalization

A Tape Catalog contains the contents of a specific tape Volume. In XenData documentation, it is termed either a 'Tape Contents Catalog' or a 'Tape Catalog' for short. The ability to create a Tape Catalog was introduced in Version 4.50 of XenData Archive Series software.

When a Volume is first added to a Volume Set, a Tape Catalog file is created for it and this is written to the disk under XenData control. As files are written to the tape Volume, the Tape Catalog is kept updated.

When a tape cartridge formatted with the TAR format becomes nearly full, a Finalization operation is automatically performed by the system. This writes the Tape Catalog to the tape and prevents more files being written it. (Finalization may also be manually initiated for a partially written tape using the XenData Management Console.) Following Finalization, the Tape Catalog is present in two locations: on the disk under XenData control and on the tape cartridge.

When a tape cartridge formatted with the LTFS format becomes full, a Finalization operation does not occur. This is unnecessary as the contents of the cartridge are maintained in the index partition of the tape.

The presence of a Tape Catalog is not required to write, read or access files from tape, but a Tape Catalog must be present on disk for successful operation of some tape management functions including Repack and to generate a report of tape or Volume contents.

When a tape formatted using TAR and containing a Tape Catalog is moved from one system to another, the Tape Catalog will automatically be copied to the disk cache of the new system, but the files on that tape will not be present in the file system interface.

When a tape formatted using LTFS is moved from one system to another, a Tape Catalog will be automatically created from the file system metadata within the index partition of the tape. The Tape Catalog will be stored on the disk cache of the new system but the files on that tape will not be present in the file system interface.

The following operations should be performed when importing a previously written tape into another system. These operations are performed using the XenData Management Console:

1. Import the tape cartridge into the tape library. This will be automatically recognized by the system and will be shown in the XenData Management Console. If the tape has been formatted with TAR, Finalized and contains a Tape Catalog, the Tape Catalog will be copied automatically to the magnetic disk volume under XenData control. However, if the tape does not contain a Tape Catalog, a Tape Catalog should first be built on the magnetic disk using the "Build Catalog" operation. This operation may take several hours because the entire tape must be scanned. If the tape has been formatted with LTFS, a Tape Catalog will be automatically created and stored on the magnetic disk volume under XenData control.
2. Next, the Import Folder Structure or Import Data operation should be performed. This publishes the current files to the file system interface. The Import Data operation not only imports the folder and file structure to the file system, it also transfers files to the disk cache according to the flush after write policy.

## 2.11 Recovering Tape Space using Repack

Repack is an operation initiated using the XenData Management Console which recovers tape space used by deleted files and old versions of files. The operation may be performed only on Volumes that are not writable, such as those that are full, finalized or write-protected. Repack is not available for WORM tapes. The repack operation performs the following:

- Files that are currently accessible via the Windows file system are copied from the selected Volume. Deleted files and old versions of files are not copied.
- These files are copied to destinations defined by the current File Group rules. A File Group rule must exist for all files that are stored on the Volume.
- After all files on the Volume have been successfully repacked, the repacked tapes are moved to the quarantined media set.

If the File Group rules have not changed since the files were first written to the repacked Volume, they will be repacked to the same Volume Set.

The repack operation should be performed on an archive system with a tape library or at least two stand-alone tape drives. The repack operation cannot be performed on an archive with only one stand-alone tape drive.

## Chapter 3. Administering the System

### 3.1 Overview of Implementing File Management Policies

The XenData Management Console is used to configure all File Group, Volume Set and tape replication options, and to view diagnostic information about the system. The XenData Management Console is a Microsoft Management Console snap-in.


All files written to the logical drive letter that is under XenData control are allocated to File Groups. Each File Group has a policy associated with it. If the File Group is written to tape, the policy is fully defined by settings made in the File Group Configuration and Volume Set Configuration sections of the XenData Management Console.

#### 3.1.1 To start the XenData Management Console

For Windows Server 2008 R2 and Windows 7:

1. Click **Start**
2. Click **Programs** (or **All Programs**)
3. Click **XenData**
4. Click **XenData System Configuration**

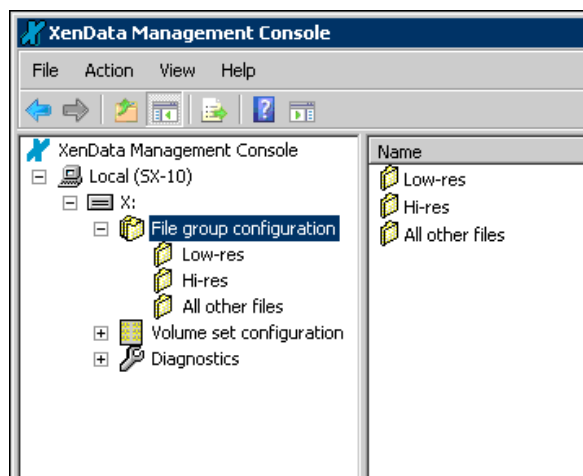
For Windows Server 2012 R2:

1. Click on the **Start** button, located in the lower-left hand corner of the screen
2. Click the arrow  near the lower-left corner of the Start screen.
3. Click on **XenData System Configuration** shown under the **XenData** category

#### 3.1.2 To navigate around the XenData Management Console

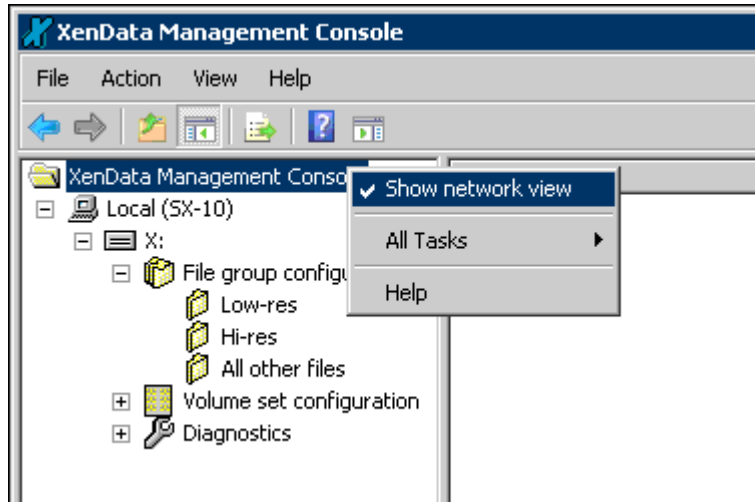
1. Click on the **+** sign next to the name of the system you wish to configure. This expands the tree structure to show the drive letter that is being controlled by the XenData6 Server software.
2. Click on the **+** sign next to the drive letter to expand it. This reveals three options: **File group configuration**, **Volume set configuration** and **Diagnostics**.
3. Click on the **+** sign next to one of the options to expand the option you wish to use.

The XenData Management Console is shown below:



### 3.1.3 To Display XenData Management Consoles for other Archives on the Network

Right-click on the text **XenData Management Console** at the top of the left pane and enable **Show network view**. The left pane then displays the XenData Management Consoles for other XenData6 Server archives on the network. In order to view other archives, the user must have appropriate security privileges.



### 3.1.4 New Installation Check List

For a new installation, after installing the XenData6 Server software and before writing any files, the Administrator must define all File Groups and Volume Sets. A useful check list of key points to be considered is given below.

- How many different File Groups are required? (Remember that a File Group is a collection of files that are all treated in the same way.)
- How many different Volume Sets are required? (Remember that a Volume Set is a set of tape cartridges that are all treated in the same way and that multiple File Groups can be allocated to the same Volume Set.)
- For each Volume Set, determine:
  - The capacity and type (WORM or rewritable) of tape cartridges
  - The format, either TAR or LTFS
  - If the 'Write to disk if no writable media are available' option is to be enabled
  - The number of replicas required, if any
  - If replication is enabled, the time schedule for generation of replicas
- For each File Group, determine:
  - If files are to be saved to tape cartridges
  - For files saved to tape, which Volume Set with its associated configuration will be used
  - If file fragmentation is to be enabled and, if so, the fragment size
  - For files archived to tape, whether the files are to be flushed from disk and if so, for what period should they be retained on disk after writing
  - For files archived to tape and scheduled to be flushed from disk, whether they will be retained on disk after reading, and if so, for what period

Care should be taken to match the File Group and Volume Settings to the capabilities of the storage hardware, especially the disk capacity, the use of a tape library and/or standalone tape drives and the number of drives in each library. It is important to match the capacity of the logical magnetic disk drive under XenData control and the file flushing policies, as there must always be sufficient capacity to store all files retained online.

Also, the number of tape drives within a library affects the number of Volume Sets that can be written simultaneously while still maintaining good performance - if significantly more tape cartridges are being simultaneously accessed than there are available tape drives, there will be excessive cartridge swaps and an associated drop in performance.

The XenData Management Console is also used for many aspects of tape cartridge management, including:

- exporting cartridges from a tape library
- finalizing tapes to write a contents catalog and prevent further writing (TAR format only)
- recording the offline location of exported cartridges
- reformat of rewritable cartridges
- repacking the contents of rewritable cartridges to release space occupied by deleted files and old file versions
- verifying data on tapes
- viewing tape cartridge information including manufacturer, serial number, etc
- viewing cartridge contents file statistics
- write protecting cartridges

## 3.2 Configuring Volume Sets

### 3.2.1 Understanding Volume Sets

A Volume Set consists of a set of tape cartridges with a defined media type (WORM or rewritable), format (TAR or LTFS) and replication strategy. Each Volume Set stores files from one or more File Groups. Automatic tape cartridge replication can be defined for each Volume Set.

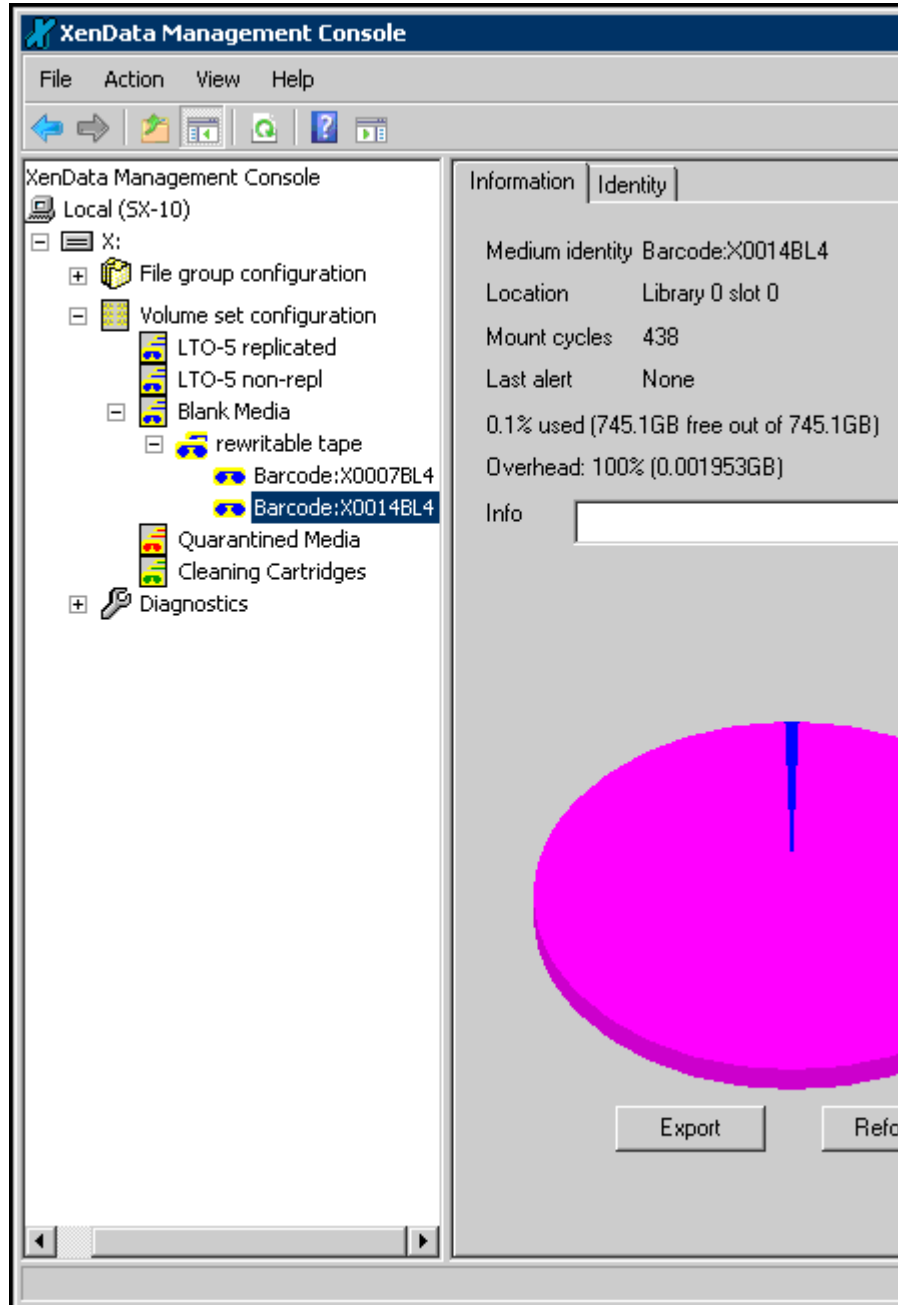
When a new system is installed, an initial Volume Set is created, ready for configuration. In addition, three special Volume Sets are shown in the XenData Management Console: the blank media set, the quarantine media set and a cleaning cartridge set.

### 3.2.2 About the Blank Media Set

The blank media set is a special Volume Set that contains tape cartridges that have been imported into a library or inserted into a drive but are not allocated to an operational Volume Set. These may be new (unused) cartridges or rewritable cartridges that have been reformatted by the system administrator. Cartridges in the blank media set are allocated to an operational Volume Set either manually by the administrator or automatically by the XenData6 Server software when all the existing cartridges in a Volume Set are nearly full. The threshold at which tapes are added to a Volume Set is determined by the Volume Set configuration settings.

When all existing cartridges in a Volume Set become full because there are insufficient blank tapes of the correct type, the system behavior depends on the Volume Set configuration settings. If the 'Write to disk if no writable media are available' option has been enabled, the system automatically enters the Pending Write Mode and will accept more data. If the 'Write to disk if no writable media are available' option has not been enabled, the system will not accept any more data and will report "disk full" when an attempt is made to write to the Volume Set.

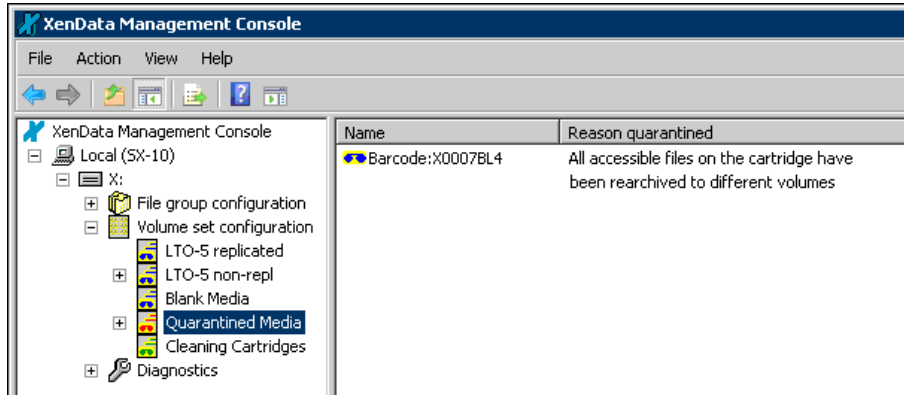




### 3.2.3 About the Quarantined Media Set


The quarantined media set is a special Volume Set that contains tape cartridges that have been imported into a library or inserted into a tape drive but for some reason cannot currently be used by the system. Typically, this will be because a cartridge has previously been used by a different application (such as a backup application), because the cartridge contents have been repacked or because an error occurred while the system was trying to identify the contents of the cartridge.

Quarantined tape cartridges must be reformatted before they can be used by the system.



### 3.2.4 About Volume Set Defaults for a New Installation

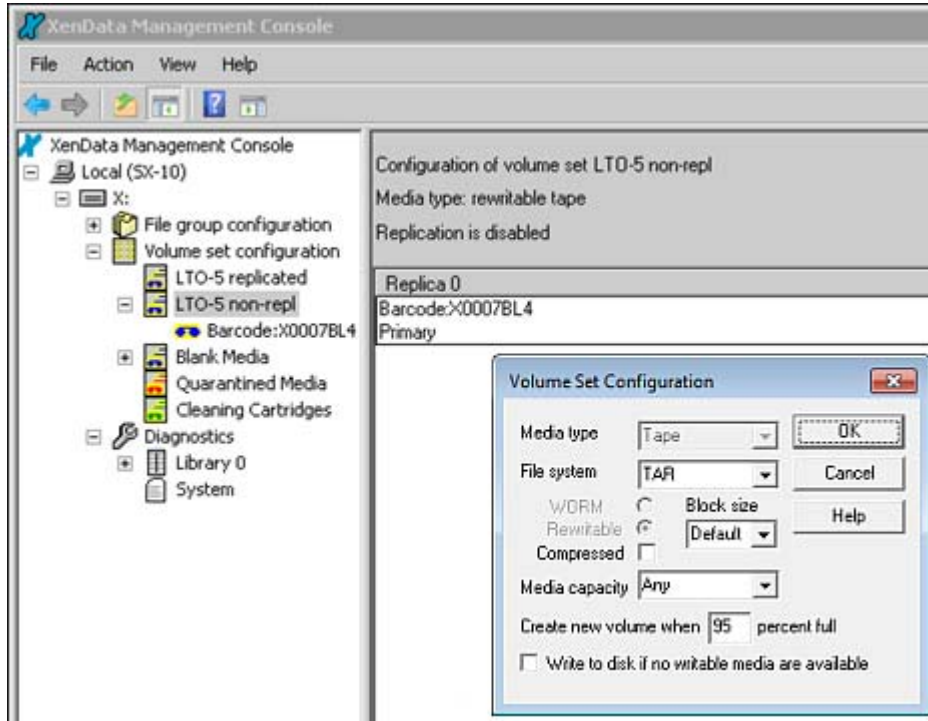
After a new installation without any tape cartridges in the library or stand-alone tape drive, the XenData Management Console will look as shown below:



With a new installation, the default system will have an initial default Volume Set, a blank media set, a quarantined media set and cleaning cartridge set. The default Volume Set will have a predefined name comprising a system unique number and numerator (e.g. 3D90284C-00000000). The default name can be edited by highlighting and selecting 'rename' via the right mouse button. The default Volume Set is setup by selecting it in the left pane and then using the 'configure' and 'replication' buttons in the right pane, as described below.

### 3.2.5 To Configure a Volume Set

1. Open the **XenData Management Console**.
2. Navigate to the **Volume Set Configuration** section.
3. Click on the required Volume Set to reveal the "Configuration of volume set" panel in the right pane of the window.
4. Click on the Configure button in the right pane.



In the Volume Set Configuration, select options as described below.

Select WORM or Rewritable tape cartridges using the two radio buttons.

Select the format, either TAR or LTFS.

The "Compressed" check box enables data compression in the tape drive. If the system is storing uncompressed data then selecting this option will allow an increased amount of data to be saved per tape cartridge. However, many applications perform their own data compression and if this is the case then it is unlikely that the compression built into the tape drive will offer any further compression. In fact, the use of tape drive compression may increase the file size for already compressed files.

Select the block size for Volume Sets that use the TAR format. Normally the default should be selected as this will select a block size that is optimized for archive and restore operations for the installed tape drive type.

You can also define a tape cartridge capacity and change the point when additional tape cartridges will be automatically allocated to this Volume Set from the Blank Media Set.

The "Write to disk if no writable media are available" check box determines the system response if all tapes in a Volume Set become full or unavailable, perhaps due to a hardware failure. If the option has been enabled by checking this box and all tapes are full or unavailable, the system automatically enters the Pending Write Mode and will accept more data which is written to the disk cache. However, if the option has not been enabled, the system will not accept any more data and will report "disk full" when an attempt is made to write to the Volume Set.

When you are satisfied with the configuration options, click the "OK" button. Clicking on "Cancel" will discard any changes you have made.


If replication is required, configure the replication parameters as described below.

### 3.2.6 To Define Replication for a Volume Set

Replication settings for a Volume Set should be defined before manually adding the first Volume. After adding a Volume, with its associated allocation of tape cartridges, only the replication schedule can be changed, not the number of replicas. Note that replication is not possible on a system that consists of only one stand-alone tape drive.

Replication is defined for a Volume Set as follows.

1. Open the XenData Management Console.
2. Navigate to the **Volume set configuration** section.
3. Click on the required Volume Set to reveal the "Configuration of volume set" panel in the right pane of the window.
4. Click on the **Replication** button in the right pane. This reveals the "Replication Configuration" dialog, as shown below.



To enable replication for the Volume Set, check the "Enable replication" box and enter the required number of replicas in the "Number of additional copies of each cartridge" box. Then define replication timing.

### 3.2.7 To Define Replication Timing

An essential part of defining the replication strategy for a Volume Set is defining the timing of replica updates. Replication schedules can specify that replication should take place "immediately"; periodically or at the same time each day.

#### To define an "immediate" replication schedule

1. Open the XenData Management Console.
2. Navigate to the **Volume set configuration** section.
3. Click on the required Volume Set to reveal the "Configuration of volume set" panel in the right pane of the window.
4. Click on the **Replication** button in the right pane. This reveals the "Replication configuration" dialog, as shown above.
5. Check the **Replicate immediately** box.

Note that immediate replication cannot be selected in conjunction with any other replication schedule. Immediate replication is not recommended for single drive library systems, as it will cause increased cartridge swapping, leading to degraded performance.

**To define a "periodic" replication schedule**

1. Open the "Replication Configuration" dialog, as described above.
2. Check the **Replicate every** box.
3. Enter the required replication schedule, using the **minutes** and **hours** radio buttons and the number entry box as appropriate.

**To configure replication at the same time every day**

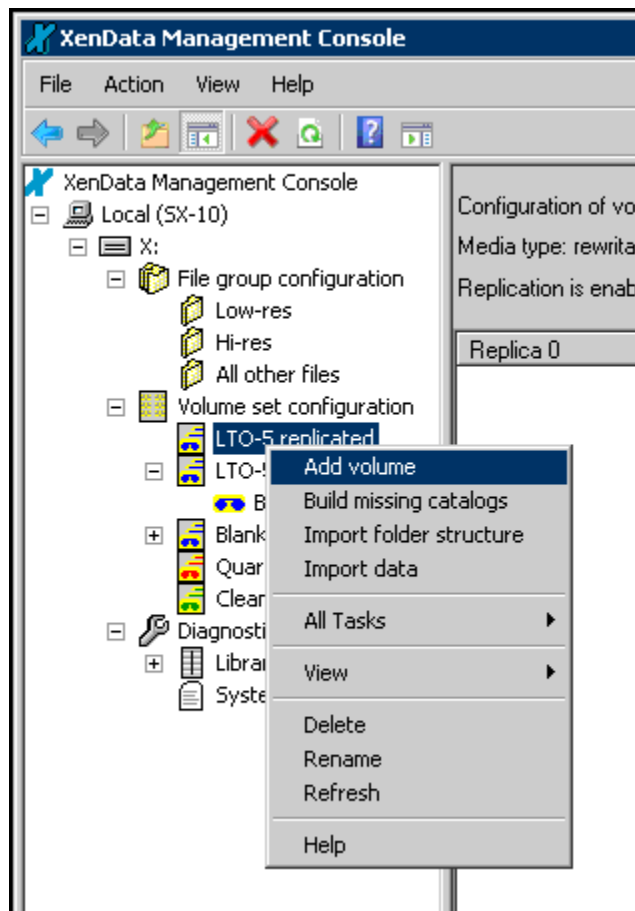
1. Open the "Replication Configuration" dialog, as described above.
2. Check the **Replicate at** box.
3. Enter the required replication time.

**3.2.8 To Allocate Cartridges to a Volume Set**

After initial creation and configuration of a new Volume Set, the system administrator must add cartridges to the Volume Set using the XenData Management Console as described below. After the initial cartridges allocated to a particular Volume Set become full, a system with a tape library will automatically add fresh cartridges from the blank media set (i.e. the administrator only needs to perform this manual operation after the initial creation and configuration of a new Volume Set). The system automatically adds blank media to the Volume Set at the percentage full point defined in the configuration for the Volume Set.

To add cartridges to a Volume Set:

1. Open the XenData Management Console.
2. Navigate to the **Volume set configuration** section.
3. Right click on the Volume Set to which media is to be added and select **Add Volume**.

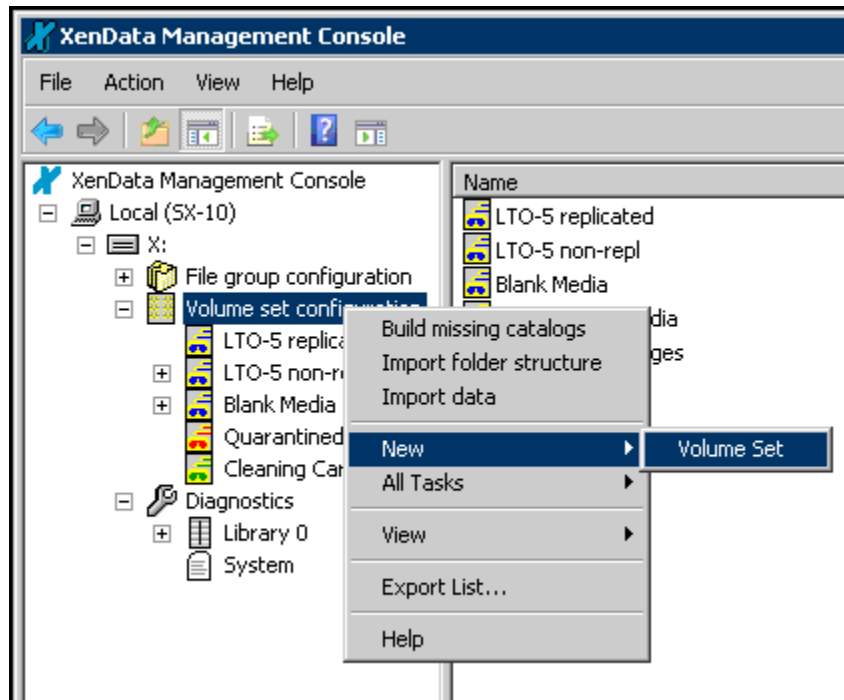


The system will then take one or more cartridges from the blank media set and allocate them to the selected Volume Set. If replication is not enabled for the selected Volume Set, only one cartridge will be taken from the blank media set. If replication is enabled, a number of cartridges will be taken from the blank media set, sufficient to satisfy the number of copies of each cartridge that was defined in the replication configuration.

**Recommendation:** Do not pre-allocate more than one tape Volume, leaving the software to automatically allocate tape Volumes as required. (Pre-allocating additional Volumes unnecessarily increases the time for the system to initialize.)

### 3.2.9 To Create a New Volume Set

1. Open the XenData Management Console.
2. Navigate to the **Volume set configuration** section.
3. Right click on **Volume set configuration**, click on **New** and select **Volume Set**.



The new Volume Set is given a unique name by the system. To rename it:

1. Right click on the new Volume Set and select **Rename**
2. Type the new name
3. Press the **Enter** key.


The new Volume Set should then be edited as described above.

### 3.2.10 To Replace a Missing Replica Cartridge

If Replication is enabled for a Volume Set, files are archived to two or more replica cartridges. If a replica tape cartridge becomes lost or damaged, it can be replaced using the **Add missing replica** operation. However, before this operation can be used, the tape that is to be replaced must be exported from the tape library and the system must be instructed to 'forget about' the missing replica.


## To Forget a Cartridge

1. Open the XenData Management Console
2. Select the tape in the XenData Management Console
3. Ensure that the tape is offline.
4. Right click on the tape and select **Forget this cartridge**



### To Add a Missing Replica

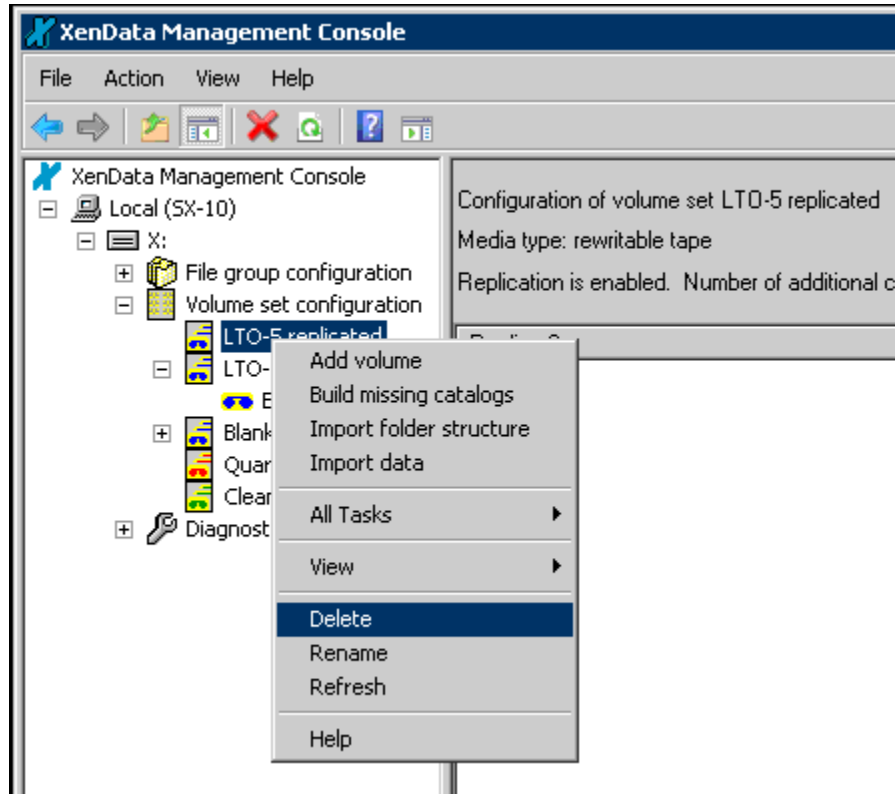
1. Open the XenData Management Console
2. Select the Volume in the XenData Management Console
3. Right click on the Volume and select **Add missing replica**





### 3.2.11 To Delete a Volume Set

1. Open the XenData Management Console.
2. Navigate to the **Volume set configuration** section.
3. Right click on the Volume Set that is to be deleted and select **Delete**.



**Note:** A Volume Set can only be deleted if it contains no tape cartridges *and* it has not been selected as the "Save to Tape" target of any File Group.

## 3.3 Managing Tapes

### 3.3.1 About Rewritable Cartridges

Unlike WORM tape cartridges, data on rewritable cartridges can be erased.

### 3.3.2 About WORM Cartridges

WORM is an acronym for Write Once Read Many. After data is written to a WORM tape cartridge, it cannot be erased or altered. The Repack operation is not available for WORM cartridges.

### 3.3.3 About Importing Cartridges

Tape cartridges are imported into a tape library using the library controls, which vary from one model to another. Alternatively, a tape cartridge is manually inserted into a stand-alone tape drive. XenData6 Server software will automatically attempt to identify any tape cartridge that is imported into a library or tape drive under its control. If the tape has been previously written by XenData software, it will be assigned to its original Volume Set. If the tape is blank, it will be assigned to the blank media set. If the tape is written with an unrecognized format, it will be allocated to the quarantined media set.

If the administrator imports a tape cartridge that was written on a different XenData system and the Volume Set does not exist on this server, it will be created automatically. However, files from this tape will not be available until the Administrator right clicks on the tape in the XenData Management Console and selects either the 'Import Folder Structure' or 'Import Data' operation.


### 3.3.4 About Reformatting Cartridges

Reformatting a cartridge erases all the data that is stored on the tape and moves the cartridge into the Blank Media Set. When a cartridge is reformatted, files that are stored only on that cartridge (for example, files that have been flushed from disk) are made inaccessible. Files that were recorded on the cartridge will be shown in History Explorer as "Archived on unknown volume".

Some tape drives do not support all of the possible reformatting functions; in these cases, only the options that are supported by the drive are presented via the XenData Management Console.


### 3.3.5 To Reformat a Rewritable Cartridge

1. Select the tape in the XenData Management Console.
2. Click **Reformat**.
3. Choose a reformat option.
  - Mark as Blank will instruct the system to treat the cartridge as blank, and move it to the blank media set. However, it will not be overwritten until it is assigned to a new Volume Set. Up to this point, the data may be retrieved by exporting and then re-importing the cartridge.
  - Low-Level format will repartition the tape. This is a quick operation that is normally used when another application has formatted a tape in a legacy format. Once this operation has been performed, it is not possible to recover data from the cartridge using standard utilities. This option is not offered if the type of drive being used does not support the operation.
  - Quick Erase uses the SCSI quick erase command, which writes an end of tape mark at the beginning of the tape (thereby marking the rest of the tape blank). This operation does not physically overwrite the data on the tape; however, once this operation has been performed, it is not possible to recover data from the cartridge using standard utilities.
  - Long Erase overwrites all the data on the cartridge. This can be a time consuming operation but it provides a good degree of certainty that the data on the cartridge cannot be recovered.
  - Overwrite Entire Medium uses write commands to overwrite all the data recorded on the cartridge with a different data pattern than that used by long erase. Using this operation followed by a long erase command more thoroughly overwrites all the data on the cartridge, for very sensitive applications.



### 3.3.6 To Write Files to Tape

To write a file to tape, the file must be allocated to a File Group and the File Group must have the "Save files to tape" option selected. The File Group will be allocated to a Volume Set and the designated Volume Set must have one or more tape Volumes with sufficient space within a tape library or stand-alone tape drive under XenData control.



Files are written via the standard Windows file system interface.

### 3.3.7 To Read Files Directly from Tape

When a user reads a file using the Windows file system interface, the system will always restore the file from the fastest physical storage device. Consequently when a file is stored on both disk and tape, it will be retrieved from disk. However, files can be read directly from tape, rather than from disk, using XenData History Explorer.

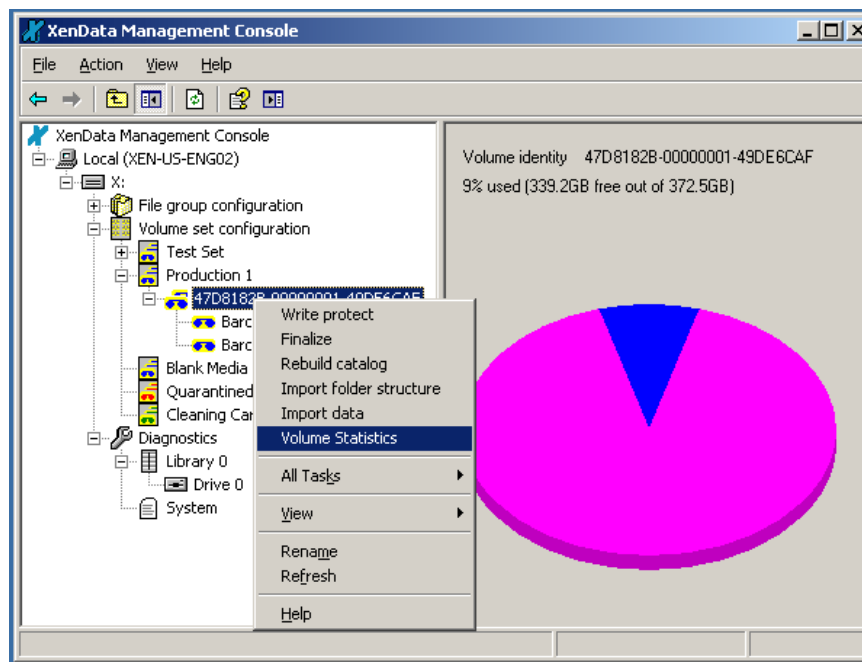
### 3.3.8. To Display File Statistics for a Volume

Remember that the XenData term 'Volume' refers to a complete set of replica cartridges which, when up-to-date, all contain the same data. If replication is not enabled, a Volume refers to an individual cartridge.

Statistics are particularly useful when considering repacking a Volume to regain space from deleted files and old file versions.

#### To Obtain Volume Statistics

1. Open the XenData Management Console.
2. Select the Volume in the XenData Management Console
3. Right click on the Volume and select **Volume Statistics**



An example of Volume Statistics is shown below.

	Number	Bytes	Bytes on media
Total number of fragment files	1896	33.34 GB	33.36 GB
Fragment files currently accessible	1890	28.65 GB	28.67 GB
Deleted fragment files	4	2.690 GB	2.690 GB
Fragment files in old versions of files	2	2.000 GB	2.000 GB
Rearchived fragment	0	0 bytes	0 bytes
Fragment files missing metadata	0	0 bytes	0 bytes
Delete and rename records	3	---	48.00 KB
Space that repack would recover			4.690 GB


The statistics provided are:

- **Total number of fragment files** This is the total number of fragment files and bytes occupied on the Volume. Note, when fragmentation is not enabled, this simply refers to the total number of files.
- **Fragment files currently accessible** This is the number of files accessible via the Windows file system and bytes occupied by them on the Volume. It excludes old versions of files and deleted files.
- **Deleted fragment files** This is the number of deleted files and the bytes occupied by them on the Volume.
- **Fragment files in old versions of files** This is the number of old (overwritten) versions of files and the bytes occupied by them on the Volume.
- **Rearchived fragments** This refers to fragment files that have been transferred to other Volumes using a repack operation.
- **Fragment files missing metadata** This refers to files that are stored on the Volume but do not have complete file system metadata stored on the disk cache.
- **Delete and rename records** These are records that are written to tape when a file is deleted or renamed.
- **Space that repack would recover** This is the space that the repack operation would recover. It is the sum of the space occupied by deleted files, old versions of files and the delete and rename records.

Note that additional information can be obtained about a Volume using the XenData Report Generator. This utility includes a report that lists all files stored on a Volume. This list may be filtered in a number of different ways and may be exported for use in Excel and other programs.

### 3.3.9 To Display Information About a Cartridge

1. Open the XenData Management Console.
2. Navigate to the **Volume Set configuration** section.
3. Click on the **+** sign next to the appropriate Volume Set to expand it and show the cartridges that are allocated to the Volume Set.
4. Click on a cartridge to display the Medium Properties pane, as shown below.




The blue segment represents used space, the pink represents free space and the red represents overhead. Overhead represents consumed tape capacity that is not used for data. If the overhead percentage is large, this is likely due to one of the following:

- only a small amount of data has been written to a tape - in this case, the overhead percentage is inaccurate
- abnormally high re-writing of data due to the drive error correction - this is indicative of a faulty drive, a drive that needs cleaning or a faulty tape
- slow writing of mainly small files or infrequent writing of mainly small files

Deletion of files in the archive file system does not cause an increase in tape overhead.

Additional information is obtained by clicking the **Identity** tab as shown below.



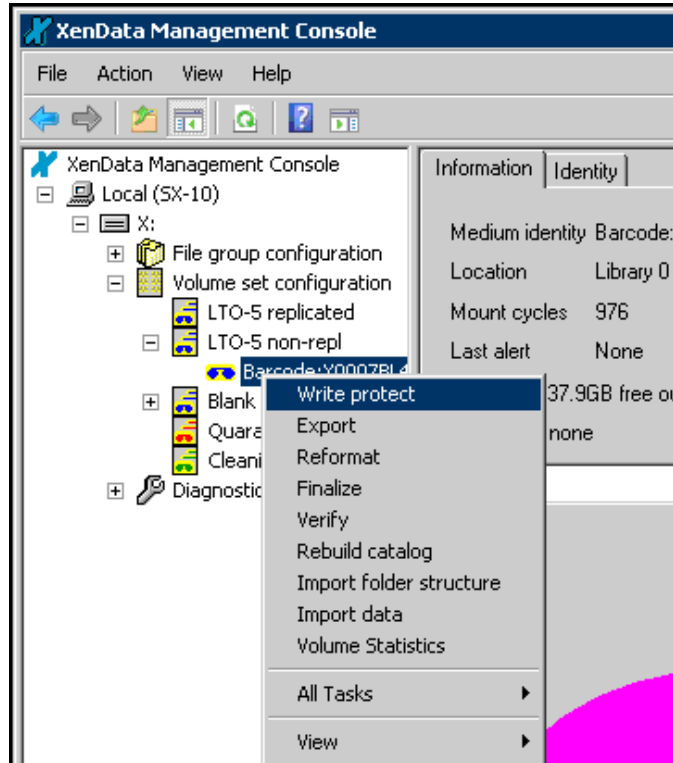
### 3.3.10 About Write-Protecting Cartridges

There may be circumstances in which a system administrator wishes to stop the system from writing data to a particular cartridge (or Volume) before the cartridge is full.

This can be achieved by write protecting the cartridge (or replicated set of cartridges representing a Volume) using the XenData Management Console. If all the cartridges in a Volume Set are write-protected, the administrator will have to add a new Volume before more data can be written to the Volume Set.

### 3.3.11 To Write-Protect a Cartridge using the XenData Management Console

1. Open the XenData Management Console.
2. Select the cartridge in the XenData Management Console.
3. Right click on the cartridge and select **Write protect**



In addition to the software write-protection described above, a cartridge may be write protected by using the write protection switch on the physical tape cartridge.

### 3.3.12 About Building and Rebuilding Tape Catalogs

A Tape Catalog contains an index of the contents of a specific Volume. The ability to create a Tape Catalog was introduced in Version 4.50 of XenData Archive Series software.

When a new Volume is first created, a Tape Catalog is created for it and written to the magnetic disk volume under XenData control. As files are written to the tape, the Tape Catalog is updated. When a TAR formatted tape is Finalized, either manually or automatically (when a tape becomes full), a copy of the Tape Catalog is written to the tape.

The presence of a Tape Catalog for a tape is required for some operations including contents repack and to view tape contents using the Report Generator. It also greatly reduces the time to perform an Import Folder Structure operation. This is an advantage when transferring tape cartridges from one XenData archive to another and when rebuilding the archive file system on the magnetic disk volume.

Where a Tape Catalog does not exist on the magnetic disk volume, one may be built using the Build Catalog operation. Additionally, a Tape Catalog may be rebuilt using the Rebuild Catalog operation. The Rebuild Catalog operation is useful if the Tape Catalog becomes corrupted perhaps because of a hardware failure.

### 3.3.13 To Build a Tape Catalog

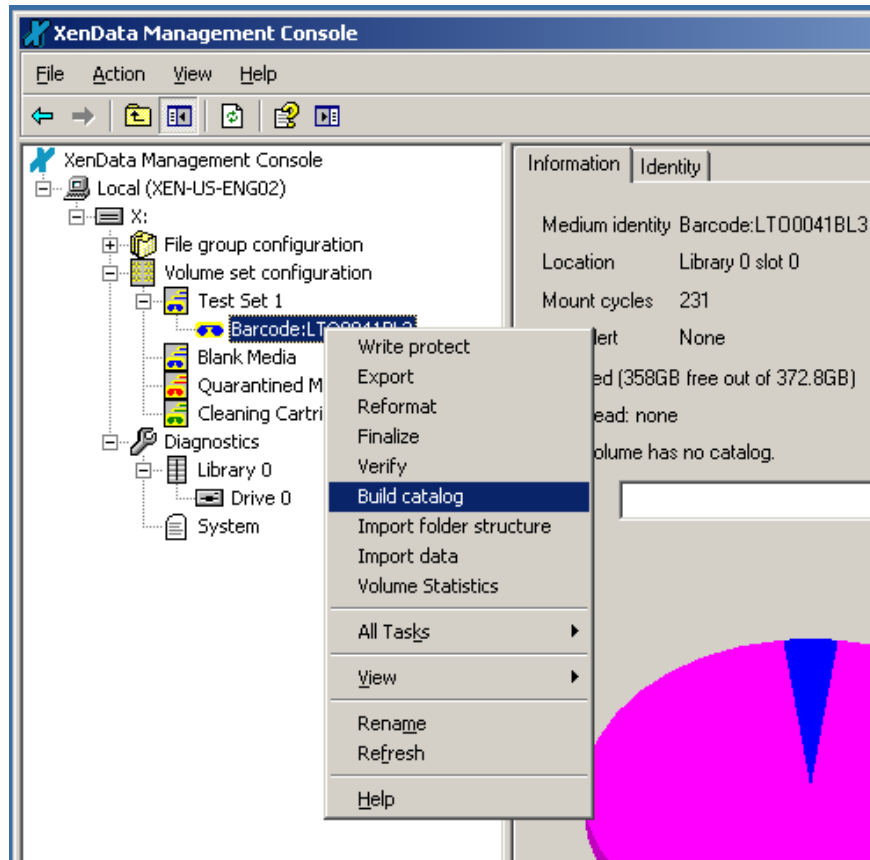
Under some circumstances, a tape cartridge may not have a Tape Catalog on the archive's magnetic disk volume. This will occur when a non-finalized tape formatted with TAR is transferred from another archive system. It also occurs with tapes that were originally added to a Volume Set using a XenData Archive Series version prior to 4.50, which is when Tape Catalogs were first introduced.



When a Tape Catalog does not exist on the magnetic disk under XenData control, the system provides a mechanism to build it, as follows:

**To Build the Tape Catalog for a Tape Cartridge**

1. Open the XenData Management Console.
2. Select the tape or Volume in the XenData Management Console
3. Right click on the tape or Volume and select **Build Catalog**.



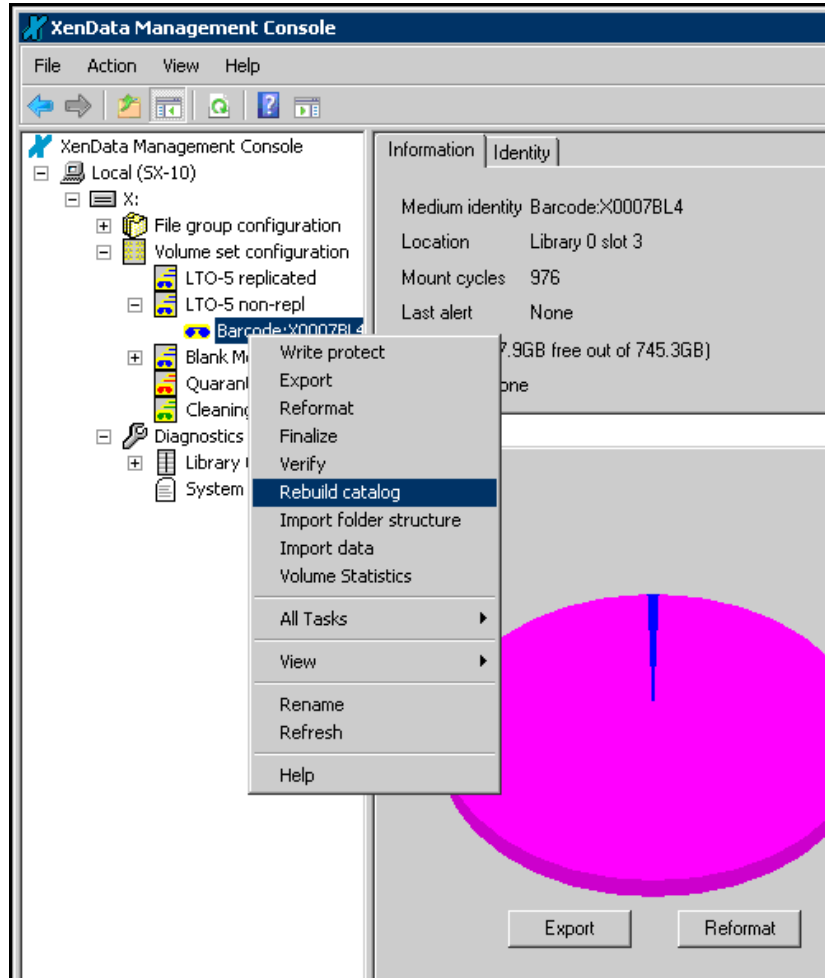
Note that the Build Catalog option is only available when a catalog is not present on magnetic disk. If a Tape Catalog already exists for the tape cartridge, the Rebuild Catalog option will be available.

**3.3.14 To Rebuild a Tape Catalog**

Under some system failure modes, the Tape Catalog stored on disk may become corrupted. If this occurs, the Tape Catalog should be rebuilt. Before performing the rebuild operation, ensure that at least one cartridge from the Volume is in a tape library or stand-alone tape drive under control of the system. The rebuild operation is performed as follows:

**To Rebuild the Tape Catalog for a Tape Cartridge**

1. Open the XenData Management Console
2. Select the cartridge or the applicable Volume in the XenData Management Console
3. Right click on the cartridge or Volume and select **Rebuild Catalog**.



### 3.3.15 To Repack a Volume

Repack is an operation which recovers tape space used by deleted files and old versions of files. The operation may be performed only on Volumes that are not writable, such as full, finalized and write-protected Volumes. The repack operation performs the following:

- Files that are currently accessible via the Windows file system are copied from the selected Volume. Deleted files and old versions of files are not copied.
- These files are copied to source destinations defined by the current File Group rules. A File Group rule must exist for every file that is stored on the Volume.
- After all files on the Volume have been successfully repacked, the repacked tapes are moved to the Quarantined Media set.


If the File Group rules have not changed since the files were first written to the repacked Volume, they will be repacked to the same Volume Set.

#### Hardware Requirements

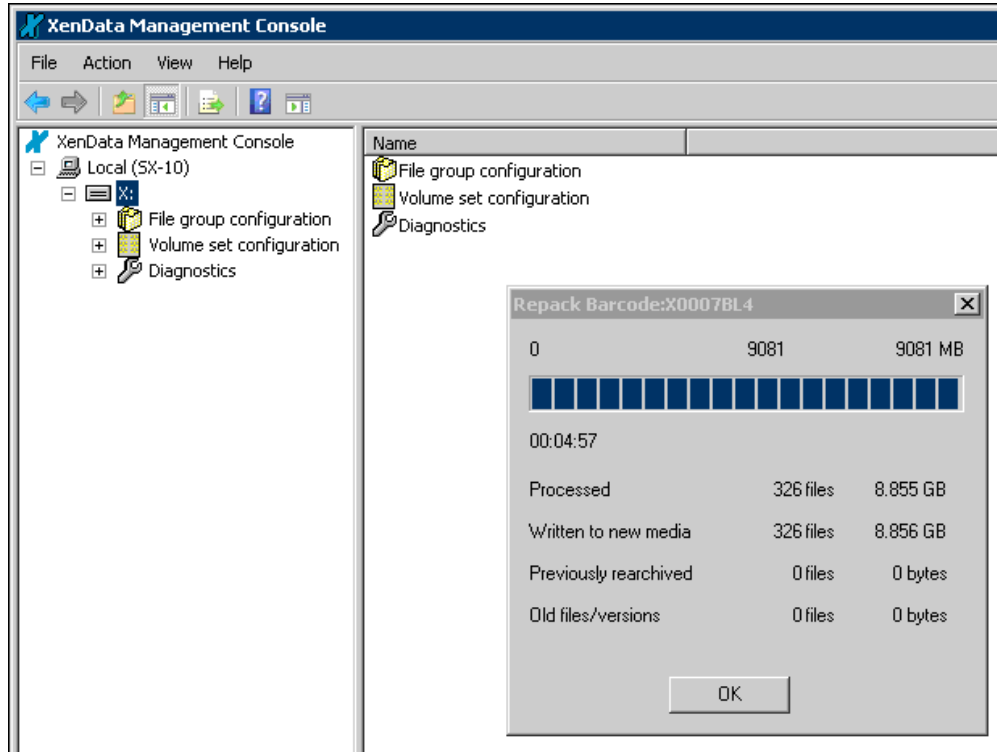
The repack operation cannot be performed on an archive with only one stand-alone tape drive.

### To Repack a Volume

1. Open the XenData Management Console
2. Select the Volume in the XenData Management Console
3. Ensure that the Volume is not writable. If it is writable, write-protect it.
4. Right click on the tape and select **Repack**

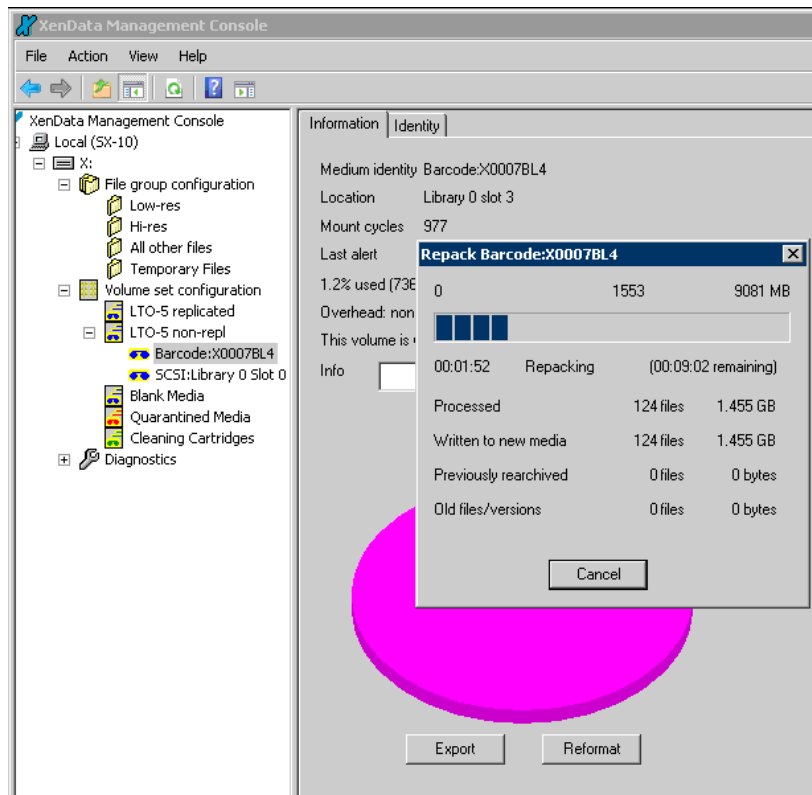


After the repack operation has been successfully completed, the tapes from the repacked Volume will be moved to the Quarantined Media set and status information as shown below will be displayed. Quarantined rewritable cartridges may then be reformatted for re-use.



### 3.3.16 To Cancel a Repack Operation

The repack operation, described above, may take many hours to complete. The operation can be cancelled and restarted at a later time. When the repack operation is occurring, a progress box is displayed as shown below and the operation can be canceled by clicking **Cancel**.



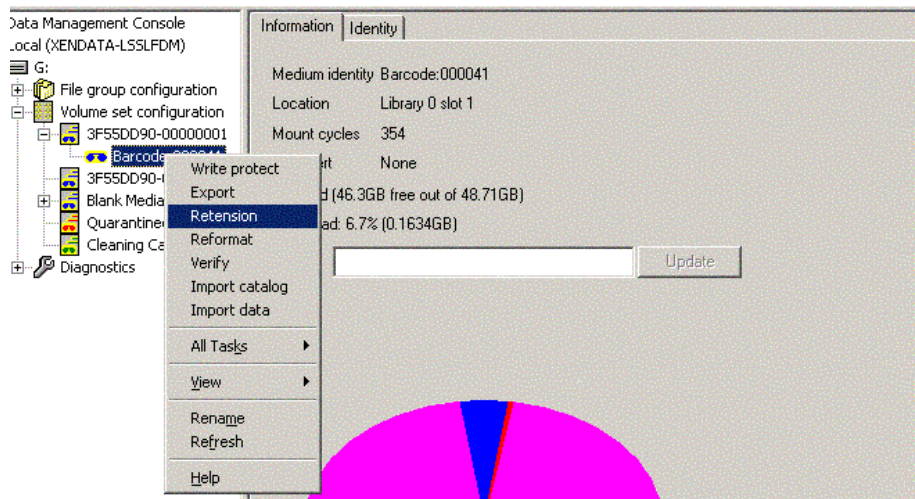
If resumed at a later time, the repack operation will restart from where it was previously cancelled.

### 3.3.17 About Re-tensioning Tapes

Under normal circumstances it should not be necessary to re-tension tapes. In exceptional circumstances, the administrator can manually re-tension a tape cartridge. Some drive types do not support the re-tension operation; in these cases, the option is not available in the XenData Management Console.

### 3.3.18 To Re-tension a Tape

1. Open the XenData Management Console.
2. Select the tape in the XenData Management Console.
3. Right click on the tape cartridge and select **Re-tension**




### 3.3.19 About Verifying the Data on a Tape

The verify function allows the administrator to check that the data on a tape can still be read, without transferring it to disk (it performs media verification). System administrators may wish to use this function to check the integrity of their archived data.

### 3.3.20 To Verify Data on a Tape

1. Open the XenData Management Console.
2. Select the tape in the XenData Management Console.
3. Right click on the tape and select **Verify**



### 3.3.21 About Finalizing a Tape

The Finalization operation is only available for tape cartridges formatted using TAR. Finalizing a tape cartridge results in special tape marks being written to the tape that prevent additional files being written to that tape. For Volumes that have a Tape Catalog on magnetic disk, this also results in a Tape Contents Catalog being written to the tape.


When a tape becomes 98% full, the system automatically finalizes it. The Administrator may finalize a partially filled tape using the XenData Management Console as described in the next section.

The presence of a Tape Contents Catalog on a tape greatly reduces the time to perform an Import Catalog operation. This is an advantage when transferring TAR formatted tape cartridges from one XenData archive to another and when rebuilding the archive file system on magnetic disk cache.

### 3.3.22 To Finalize a Tape

When a TAR formatted tape becomes 98% full, the system automatically finalizes it. Tapes may be finalized before they become 98% full as described below.


1. Open the XenData Management Console.
2. Select the tape in the Management Console.
3. Right click on the tape and select **Finalize**.
4. Press **Enter**.



### 3.3.23 To Export a Cartridge

Cartridges should be exported from a tape library using the XenData Management Console, as described below:

1. Open the XenData Management Console.
2. Navigate to the **Volume set configuration** section.
3. Click on the **+** sign next to the appropriate Volume Set to expand it and show the tapes that are allocated to the Volume Set.
4. Click on a tape to display the Medium Properties pane, as shown below.
5. Click **Export** in the right pane of the console or right-click the tape in the left pane and select **Export**.



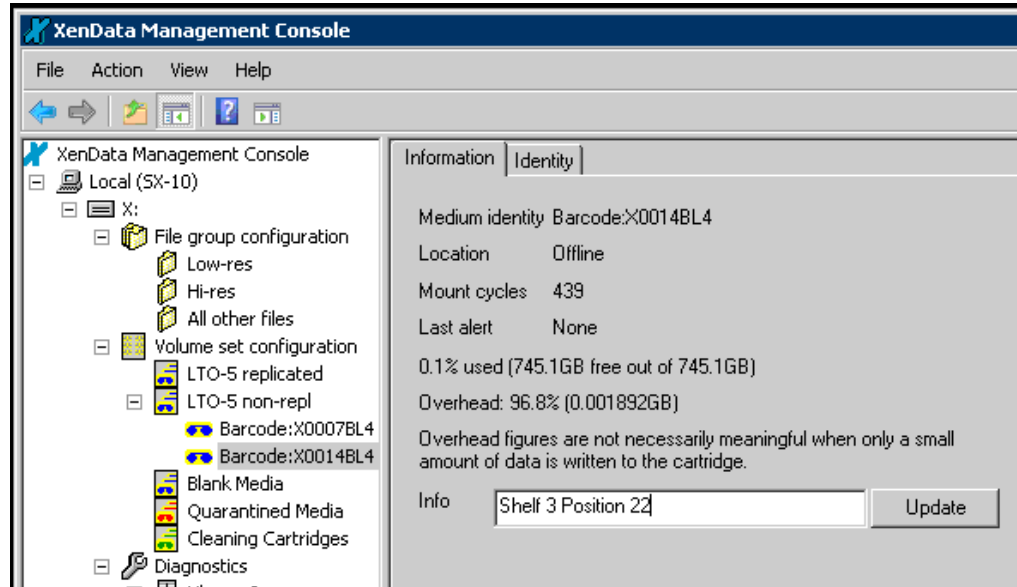
The selected tape will be moved to one of the tape library mail slots (often called I/O slots or I/E Elements) or ejected from a stand-alone drive.

When using a stand-alone tape drive, a tape cartridge may be exported (ejected) either by using the procedure described above or by pressing the tape drive eject button.

### 3.3.24 To Set the Administrator Defined Information for a Cartridge

1. Open the XenData Management Console.
2. Navigate to the **Volume set configuration** section.
3. Click on the **+** sign next to the appropriate Volume Set to expand it and show the tapes that are allocated to the Volume Set.
4. Click on a tape to display the Medium Properties pane, as shown below
5. Enter the desired information in the "Info" box
6. Click on "Update" to save your changes.





## 3.4 Allocating Files to File Groups

### 3.4.1 About File Groups


A File Group is a collection of files that have the same file management policy and consequently are all treated in the same way by the system.

Whenever a file is written to the archive, XenData6 Server software needs to know how to treat it. The way the file should be handled is defined by File Group rules, so the first thing the system does when a file is opened or created is to allocate it to a File Group.

Every file that is handled by the system belongs to exactly one File Group. Files are assigned to a File Group on the basis of their name and path. This assignment can be based on the name of the folder that contains a file, the name of the file or both.

### 3.4.2 To Allocate Files to a File Group

1. Open the XenData Management Console.
2. Navigate to the **File group configuration** section.
3. Click on the required File Group to reveal the "Configuration of file group" panel in the right pane of the window.
4. Update the **File name or path pattern** box with the appropriate text, as described below
5. If required, update the **Exclude pattern** box, as described below.
6. Click **Apply**.



### 3.4.3 To Construct the "File Name or Path Pattern" or the "Exclude pattern"

The system maintains a list of File Groups, each of which has a "File name or path pattern". To establish which File Group to use for a particular file, the system starts at the top of the list of File Groups and tries to match the file name to the pattern for the File Group. If the file name matches the pattern for the first File Group then the file will be allocated to the first File Group. If the file name does not match the first path pattern, the system tries the next File Group in the list, and so on down the list until it finds a match or reaches the end of the list. If it reaches the end of the list, the system blocks opening or creation of the file (it returns an error to the application that tried to open the file).

Files are allocated to File Groups based on their folder name, file name, filename extension or a combination of these. Standard filename and wildcard conventions (such as "\*" and "?") may be used during the pattern match. As an extension to normal pattern matching syntax, the special directory wild card '...' can be used to match any number of intermediate sub-folders. The system supports multiple patterns per File Group, separated by semicolons.

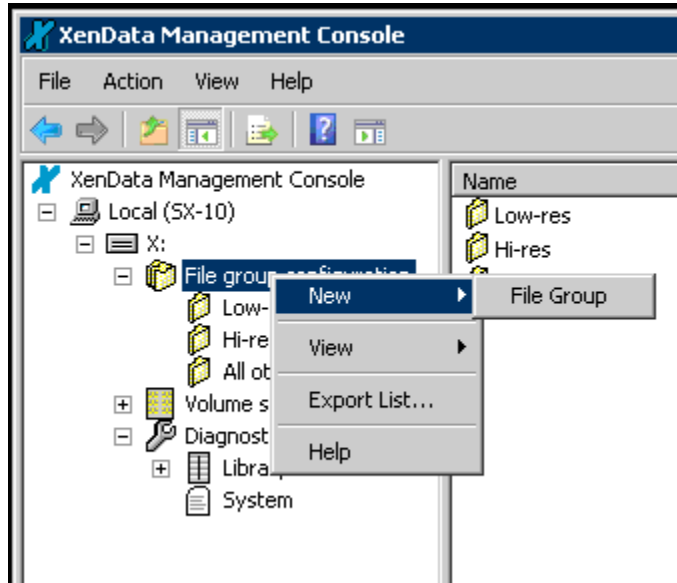
Some example file name or path patterns are:

- a. \*.tif selects files with the extension .tif for the File Group.
- b. abc???.tif selects files that start with abc, have the extension .tif and have six characters before the extension.
- c. \Images\\* selects files that are in the folder \Images.
- d. \Images\...\\* selects files that are in the folder \Images or any of its sub-folders.
- e. \Images\...\\*.tif selects files with the extension .tif that are in the folder \Images or any of its sub-folders.

### 3.4.4 To Create a New File Group

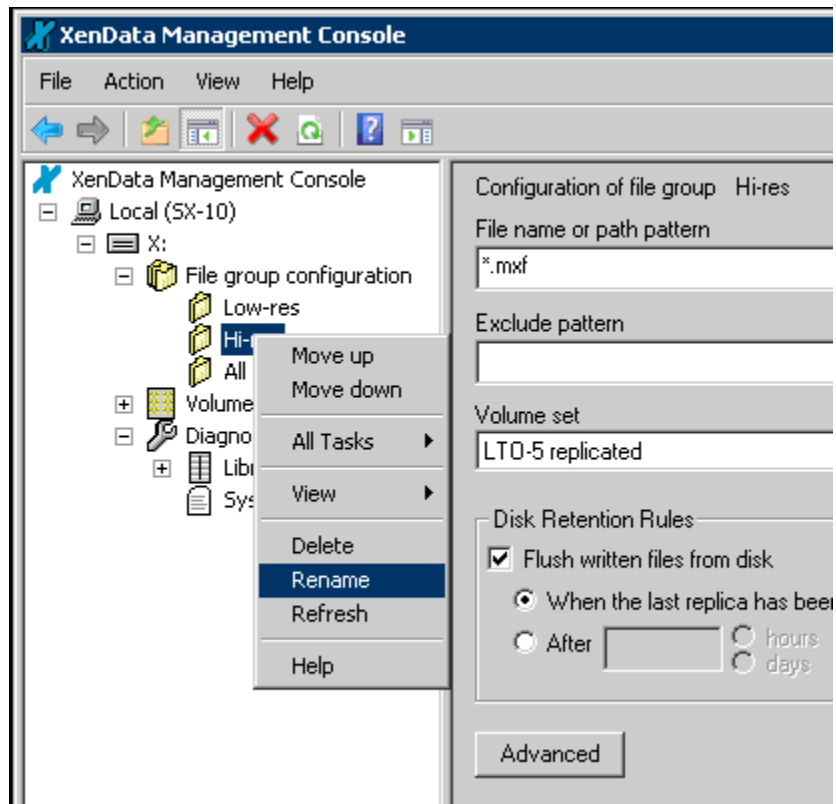
1. Open the XenData Management Console.
2. Navigate to the **File group configuration** section.
3. Right click on **File group configuration**, click on **New** and select **File Group**.

The new File Group will be named <<new>> and should be renamed as described below.



### 3.4.5 To Rename a File Group

1. Open the XenData Management Console.
2. Navigate to the **File group configuration** section.
3. Right click on the File Group that is to be renamed and select **Rename**.
4. Type a new name for the File Group
5. Press **Enter**.




### 3.4.6 About Changing the Order of File Groups

The order of File Groups in the XenData Management Console is important because an individual file can be allocated to only one File Group and the allocation rules are applied in the order that the File Groups appear in the left pane of the console with files being allocated to the uppermost applicable File Group.

### 3.4.7 To Change the Position of a File Group in the List

1. Open the XenData Management Console.
2. Navigate to the **File group configuration** section.
3. Right click on the File Group that is to be moved and select **Move up** or **Move down**.



### 3.4.8 Examples of Allocating Files to File Groups

The easiest way to illustrate how to allocate files to File Groups is by way of examples and a number of these are given below. In each case, files are allocated to three different File Groups.

**Example 1:** One File Group contains all files with names ending in “.tif”; a second File Group is for all files with names ending in “.txt”; and a third File Group contains all other files.

File Group 1	File name or path pattern:	<b>*.tif</b>
	Exclude pattern:	
File Group 2	File name or path pattern:	<b>*.txt</b>
	Exclude pattern:	
File Group 3	File name or path pattern:	<b>*</b>
	Exclude pattern:	

In this example, no path has been specified and consequently the file name rules apply to all files written to the logical drive managed by XenData6 Server software, no matter to which folder they are written.

Note that the “Exclude pattern” boxes are empty in this example. Note also that we used “\*” rather than “\*.\*” in File Group 3 to ensure that all files are included in the File Group including those without a name extension.

**Example 2:** One File Group contains all files written to a folder at the root called “\project01\”; another contains all files written to a folder at the root called “\project02\”; and a third File Group contains all other files.

File Group 1	File name or path pattern:	<b>\project01\*</b>
	Exclude pattern:	
File Group 2	File name or path pattern:	<b>\project02\*</b>
	Exclude pattern:	
File Group 3	File name or path pattern:	<b>*</b>
	Exclude pattern:	

Note that “Exclude pattern” boxes are empty in this example. Note also that we used “\*” rather than “\*.\*” in to ensure that all files are included in the file group including those without a name extension.

**Example 3:** This is similar to example 2, but includes all sub-folders for the first two File Groups. One File Group contains all files written to a folder at the root called “\project01\” and its sub-folders; another contains all files written to a folder at the root called “\project02\” and its sub-folders; and a third File Group contains all other files.

File Group 1	File name or path pattern:	<b>\project01\...\*</b>
	Exclude pattern:	
File Group 2	File name or path pattern:	<b>\project02\...\*</b>
	Exclude pattern:	
File Group 3	File name or path pattern:	<b>*</b>
	Exclude pattern:	

In this example, the use of “...\\*” denotes the specified path and all folders below it.

**Example 4:** This is similar to example 3, but all temporary files are excluded from the first two File Groups by using the “Exclude pattern” box. Consequently, all file names ending in “.tmp” are allocated to the third File Group.

File Group 1	File name or path pattern:	<b>\project01\...\*</b>
	Exclude pattern:	<b>*.tmp</b>
File Group 2	File name or path pattern:	<b>\project02\...\*</b>
	Exclude pattern:	<b>*.tmp</b>
File Group 3	File name or path pattern:	<b>*</b>
	Exclude pattern:	

**Example 5:** This is similar to example 4, where all temporary files are excluded from the first two File Groups by using the “Exclude pattern” box. As for example 4, all file names ending in “.tmp” are allocated to the third File Group. However, the administrator has not set up a ‘catch-all’ File Group rule at the bottom of the File Group list. In this case, the system will not allow writing of files unless they are written to the folders project01\, project02\ or their sub-folders or the files end with “.tmp”.

File Group 1	File name or path pattern:	<b>\project01\...\*</b>
	Exclude pattern:	<b>*.tmp</b>
File Group 2	File name or path pattern:	<b>\project02\...\*</b>
	Exclude pattern:	<b>*.tmp</b>
File Group 3	File name or path pattern:	<b>*.tmp</b>
	Exclude pattern:	

**Example 6:** This example illustrates the importance of File Group order.

File Group 1	File name or path pattern:	<b>\project01\*</b>
	Exclude pattern:	
File Group 2	File name or path pattern:	<b>*txt</b>
	Exclude pattern:	
File Group 3	File name or path pattern:	<b>*</b>
	Exclude pattern:	

In this example, files ending in ".txt" in folder project01\ are allocated to the same File Group as the other files in this folder. If the order of the first two rules was changed, files ending in ".txt" would be allocated to the same File Group as the ".txt" files in the other folders.

### 3.5 Setting Policies for File Groups

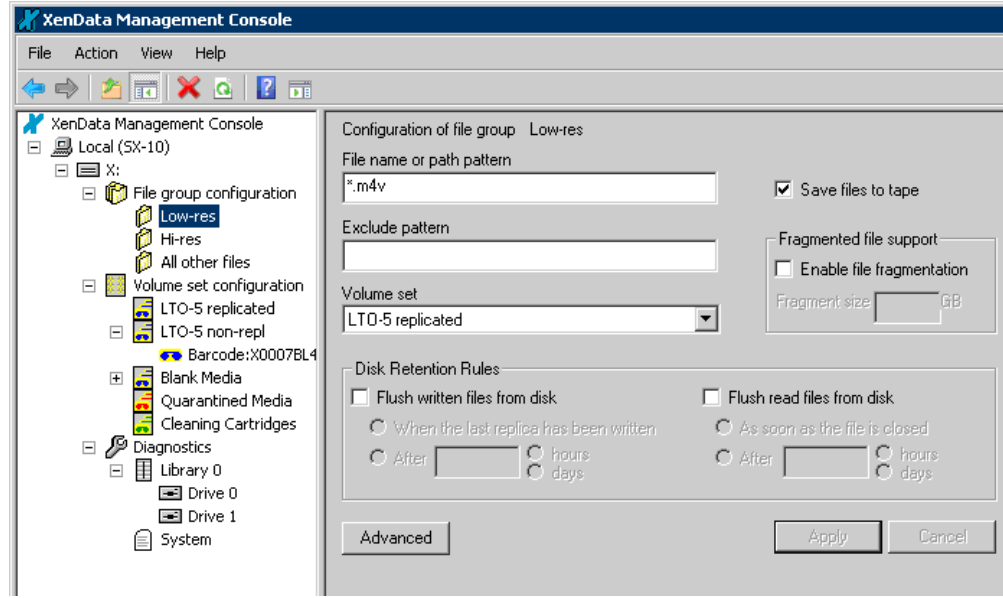
#### 3.5.1 About Storage Options for a File Group

The administrator must define storage rules for each File Group. The following policy options must be set:

- File system physical storage, including whether the File Group is stored on tape and, if so, on which Volume Set.
- Settings for File Fragmentation.
- Disk Retention Rules
- Advanced Settings

#### 3.5.2 To Select Storage Options for a File Group

1. Open the XenData Management Console.
2. Navigate to the **File group configuration** section.
3. Select the File Group that is to be configured by clicking on it.
4. Determine whether files in the File Group are to be saved to tape, and enable the "Save files to tape" option if appropriate.
5. Determine whether or not file fragmentation is required for the File Group and enable if appropriate. File fragmentation is typically required only for very large files, and is described below.
6. If files are being saved to tape, determine whether to use file flushing to save space on magnetic disk. If file flushing is to be used, configure the Disk Retention Rules as described below.
7. Click on the **Advanced** button near the bottom of the screen if Advanced Settings are needed and make additional selections.



### 3.5.3 File Fragmentation

XenData file fragmentation provides the ability to partially read files which can lead to significant performance improvements with large files. Furthermore, file fragmentation allows a large file to be spanned across multiple Volumes.

Enabling file fragmentation typically provides benefits for File Groups that predominantly consist of large multi-gigabyte files. If file fragmentation is enabled, the system has the following characteristics:

- when a portion of a file is read from tape, only the applicable fragments will be read, saving both transfer time and space on the disk cache;
- when an application modifies a large file by appending, the appended data will be written to tape as one or more additional fragments, saving space on the tape cartridge;
- if an application modifies a small part of a large file, for example by updating an index at the beginning of the file, then only the fragments containing modified data will be written to tape; and
- on writing a file, the system may span multiple Volumes if the File Group advanced option to permit file spanning is enabled and, when spanning occurs, complete fragments of spanned files will be written to each of the spanned Volumes.

If file fragmentation is not enabled, the system has the following characteristics:

- when a file or portion of a file is read from tape, the complete file will be read from tape;
- when a file is modified, the new version of the file will be completely written to tape; and
- on writing a file, the system will always write the complete file to a single Volume.

If the administrator enables file fragmentation for a File Group, the fragment size must be defined.

Recommended fragment sizes depend on the application and tape drive transfer rates but for LTO-5 will typically be 5 GB or larger.

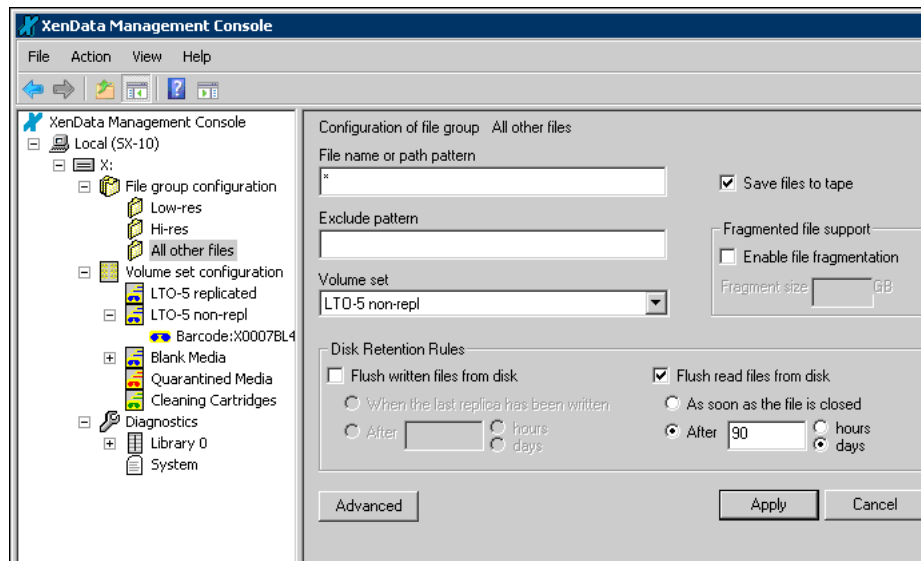
### 3.5.4 Disk Retention Rules

The administrator can configure the system such that after a file has been securely written to tape, the instance stored on disk will be flushed to release the disk space that was occupied by the file.

The flush operation causes the file to be removed from disk, but it is still visible in the file system and accessible to normal applications from tape. This functionality is enabled by configuring the Disk Retention Rules in the File Group options. The Disk Retention Rules are only available for File Groups where "Save files to tape" has been enabled.

Characteristics of flushed files are as follows:

- Flushing from disk does not affect the presence and location of a file within the file system.
- File properties - including file size - do not change, except the Windows offline attribute is set.
- Flushed files are restored from tape by simply reading the file. As long as the file is available on near-line tape, the system will seek to the beginning of the file and will restore the file automatically.



If all the boxes in the Disk Retention Rules area are unchecked, the system retains all current files within the File Group on magnetic disk indefinitely. A number of options are available for flushing current files from disk, as described below.

When a file is first written to the drive letter under XenData control, it is always first written to magnetic disk. The options, which are described below, allow the administrator to flush files from disk after completion of the write operation. Note, the system takes great care not to flush files from disk until the last replica has been securely written to tape.

When a file has been flushed from disk by setting one of the options described below, the system restores the file to disk each time that it is read from tape. The administrator can use the options described below to flush files from disk after the application that read the file has finished with it.

### Options for Disk Retention


- All current files are retained on disk indefinitely  
Ensure that all boxes in the Disk Retention Rules area are unchecked.
- Files are flushed immediately after writing  
Select **Flush written files from disk** and **When the last replica has been written**. With these options selected, files will be flushed as soon as the last replica has been completely written to tape. However, when a file is subsequently read, it will be restored to disk and retained indefinitely unless **Flush read files from disk** is checked, as described in options f. and g. below.



- c. Files are flushed a preset length of time after writing  
Select **Flush written files from disk** and **After** and put an appropriate number in the box. With these options selected, files will be flushed a defined length of time after they were written, or when the last replica has been completely written to tape, whichever comes later. However, when a file is subsequently read, it will be restored to disk and retained indefinitely unless **Flush read files from disk** is checked, as described in options h. and i. below.
- d. Files are flushed immediately after being read  
Select **Flush read files from disk** and **As soon as the file is closed**. Ensure that **Flush written files from disk** is unchecked. With these options selected, files will be flushed as soon as the file is closed. Note that files that have been written to the system will be retained on disk indefinitely, until they are read.
- e. Files are flushed a preset length of time after being read  
Select **Flush read files from disk** and **After** and put an appropriate number in the box. Ensure that **Flush written files from disk** is unchecked. This option is similar to the option d. above, but waits for the defined period before it applies its rule. As with option d., files will be retained on disk indefinitely, unless they are read.
- f. Files are flushed immediately after writing and after each subsequent read  
Select **Flush written files from disk** and **When the last replica has been written**. Also, select **Flush read files from disk** and **As soon as the file is closed**. With these options selected, files will be flushed as soon as the last replica has been completely written to tape. Each time a file is subsequently read, it will be temporarily restored to disk, but when the file is closed, it will immediately be flushed from disk again.
- g. Files are flushed immediately after writing and a preset time after each subsequent read  
Select **Flush written files from disk** and **When the last replica has been written**. Also, select **Flush read files from disk**. Then select **After** under the "Flush read files from disk" box and enter the required number of hours or days for retention of the file on disk after it is last read. With these options selected, files will be flushed as soon as the last replica has been completely written to tape. However, each time a file is subsequently read, it will be temporarily restored to disk, but will be flushed from disk again when the defined time has elapsed following the last time the file was closed.
- h. Files are flushed from disk a preset time after writing, combined with immediate flushing after being read  
Check both the **Flush written files from disk** and **Flush read files from disk** boxes. Then select **After** under the "Flush written files from disk" box and enter the required timing for the flush operation following writing. Then select **As soon as the file is closed** under the "Flush read files from disk" box. Under this set of options, following completion of replication, the system will keep files on disk until they are read or until the defined time has elapsed since they were written, whichever comes first.
- i. Files are flushed from disk a preset time after writing unless the file is read within a (possibly different) preset time  
Select both the **Flush written files from disk** and **Flush read files from disk** boxes. Select **After** under the "Flush written files from disk" box and enter the required time for the flush operation after writing. Then select **After** under the "Flush read files from disk" box and enter the required time for retention of the file on disk after it is last read. Under this set of options, following completion of replication, the system will keep files on disk until completion of the "written" or the "read" interval, whichever is later.

### 3.5.5 File Group Advanced Options

The file group advanced options dialog includes settings that can provide improved performance in certain situations, or offer enhanced data integrity at the expense of performance for some applications.



### Available Options

- a. Keep all versions on disk  
XenData6 Server software normally maintains a complete version history of every file under its control. However, only the most recent version of a file is available to standard applications through the file system interface; older versions of files will usually never be read and if required, they must be recovered using XenData History Explorer. To conserve disk space, the system removes old versions of files from disk, making them available only from tape. Selecting the "Keep all versions on disk" option changes this behavior so that all old versions of a file are retained on disk.
- b. Do not preserve history for deleted files  
In normal operation, XenData6 Server software maintains the entire version history of all files under its control. For compliance applications, this is the functionality that is required. However, certain applications produce large numbers of intermediate files that are valuable at the time they are created, but at some later time the intermediate files are no longer required and the application deletes them. Applications that display this behavior might be collecting incoming data into intermediate files and later merging them into larger container files for long-term archival storage (some email archiving systems work this way). Because the incoming files (for example, individual email messages) are valuable, they should be archived to tape, providing a backup of the disk. However, after the application has deleted them, the intermediate files are no longer needed and there is no requirement to continue to maintain their history. Maintaining a file's history consumes a small amount of space on the magnetic disk (for metadata) and if the system is maintaining metadata for a very large number of deleted files, the space consumed by their metadata may become unacceptably large. Selecting this option removes the metadata for deleted files and allows the system to recover the space that would otherwise be consumed. Files that belong to file groups with this option selected will not be visible in History Explorer after they have been deleted.
- c. Use a background process to archive old files if updated file rules indicate that files that were not previously archived should now be archived  
If a File Group rule does not have "save to tape" enabled, files will be written to the disk cache only. If the File Group rule is then changed by enabling the "save to tape" option, new files written to the archive will be archived to tape but files written prior to the change will not, unless this option, "Use a background process...", is enabled.

The “Use a background option...” causes the system to perform a background scan and archive old files to tape that should now be archived according to the changed File Group rule.

- d. Use compact metadata representation to save space by reducing the amount of information retained for the History Explorer

Part of the information shown by XenData History Explorer is the creation date for every version of a file. This information is stored in the file's metadata, and it is only used by the History Explorer (or any other application that fetches it using the XenData Archive Series API). If a very large number of versions of a file are created, this extra information in the metadata increases the size of the metadata file, consuming space on the disk cache. This option changes the format of the metadata so that file version modification dates are not saved, thereby reducing the size of the metadata. For most archiving applications, this optimization is not worthwhile because the number of file versions created is small and no space will be saved.

- e. Allow file fragments to span more than one volume

This option is applicable when file fragmentation is enabled and it determines whether or not a file may span Volumes. When enabled, the file fragments for a file may be written so as to span Volumes. When this option is not enabled, all file fragments for a particular version of a file will always be written to only one Volume.

- f. Archive file fragments as soon as they are available

In normal operation, a file's data is written to tape after it has been successfully written to magnetic disk. Some application types, for example certain backup applications, can create very large files that are written sequentially (i.e. the application starts to write data at the beginning of the file and proceeds linearly through the file, gradually extending it as more data is written). In this situation it would be more efficient to write data to tape as soon as the application has finished writing the first fragment, rather than waiting for the application to write the entire file. This can be achieved by enabling file fragmentation and selecting this option.

- g. Force the first read of a file after it is written to come from the archive, so that the integrity of the archive may be checked.

Some applications employ a read-after-write check to verify the integrity of data written to the archive. However, the normal behavior of the system is always to read data from the fastest available location. For data that has just been written to the system, this will usually be the magnetic disk cache (or even an intermediate RAM cache). Performing a read-after-write check would check the integrity of the data in the cache, rather than on the tape. This option forces data to be read from the tape, even if it is available from an intermediate cache, thereby allowing applications to verify the integrity of data on the tape. Note that only the first read is forced to come from the tape; subsequent reads will be satisfied from the cache if possible.

- h. Make files read-only

This option forces all files in the file group to be marked "read-only". This read-only attribute cannot be changed after a file has been created.

## Chapter 4. Using Mac Clients

### 4.1 About Using Mac Clients

Mac OS X clients should use the SMB/CIFS network protocol when archiving to and restoring from an archive running XenData6 Server software.

When a Mac OS X client computer writes to a Windows SMB share (such as a XenData archive), it either uses 'AppleDouble' files or alternate data streams. This version of XenData6 Server fully supports AppleDouble files and manages them using hidden File Group rules; it also supports alternate data streams when writing to LTO cartridges using the TAR format; however, it does not support alternate data streams when writing to LTO cartridges using the LTFS format.

There are three important areas of file management which are handled automatically by this version of XenData6 Server:

- AppleDouble files (consisting of a data fork file and a resource fork file) are managed by hidden File Group rules;
- Apple hidden Desktop Services Store files (.DS\_Store) are also managed by hidden File Group rules; and
- Alternate Data Streams (an alternative to AppleDouble files) are managed by the XenData software when writing to TAR format tape cartridges and hidden file Group rules manage the streams.

In addition, there are two areas of file management which require configuration of each Mac client that has access to the XenData archive:

- Icon Preview must be disabled for all Mac clients that access the XenData archive, otherwise an excessive number of unwanted file restores may result which can greatly impair the archive's performance; and
- Alternate Data Streams must be disabled for Mac clients that write and restore files to LTO cartridges formatted with LTFS.

The hidden File Group rules that automatically manage AppleDouble files, streams and .DS\_Store files are described in more detail in the next section.

### 4.2 Hidden File Group Policies

In addition to the File Group rules defined by the administrator using the XenData Management Console, hidden File Group rules are automatically implemented for correct management of files that are created by Apple Mac clients.

Desktop Services Store files (named .DS\_Store) are hidden files created by the Mac OS X Finder in every folder that it accesses. Finder uses these to retrieve custom attributes of a folder such as background color and position of icons. The XenData system will automatically store any file named .DS\_Store on the disk cache but will not save the file to tape. This rule overrides all policies defined in the XenData Management Console.

When a Mac client computer running OS X writes to a Windows SMB share, it writes either AppleDouble files or files with an alternate data stream. The handling of these two cases is described below.

An AppleDouble file consists of a data file (often called a data fork) and an associated resource fork file. The resource fork is a hidden file which is typically very small and has the same name as the data file with '.\_' (dot under-score) prepended. For example, if QuickTime Pro creates a file called 'abc.mov' and the file is saved to a Windows share, the data fork file will also be named 'abc.mov' and the hidden resource fork file will be named '.\_abc.mov'. With a data tape archive system, it is important that resource fork files are not flushed as these are accessed frequently by the Mac client. The XenData system will automatically prevent flushing of all resource fork files. In cases where the corresponding data file is saved to tape, the resource fork will also be saved to the same Volume Set. These rules override all policies defined in the XenData Management Console.

When OS X writes a file with an alternate data stream to a XenData archive, the alternate data stream is very small, as it contains the same information as held in the resource fork when AppleDouble files are written. As for resource forks, the alternate data stream is accessed frequently by Mac clients and consequently it is important that the alternate data streams are not flushed. The XenData system will automatically prevent flushing of all alternate data streams and when written to tape, the system will ensure that the main file stream and any alternate data streams are written to the same Volume Set. These rules override all policies defined in the XenData Management Console.

#### 4.3 Disabling Icon Preview

When a folder is opened using Finder on a Mac client computer, all files contained in that folder are read if 'show icon preview' is enabled. This will cause unnecessary file restores from the archive. The best way to avoid this problem is to disable this preview option within Finder. This setting is found in the Finder menu at View>Show View Options>Show icon preview.

#### 4.4 Disabling Alternate Data Streams

From OS X version 10.6, Apple uses alternate data streams as the default configuration rather than writing AppleDouble files when writing to a XenData SMB share. However, the XenData archive does not support alternate data streams when writing to LTFs formatted LTO cartridges and, in this case, the use of alternate data streams must be disabled for all clients writing to LTFs formatted tapes. The recommended method for disabling alternate data streams is described in XenData Technical Note XTN1001.

#### 4.5 Support of OS X Characters

Mac OS X supports characters within file and folder names that are invalid on Windows systems. These characters are / ? < > \ : \* | " and any character you can type with the Ctrl key. With XenData6 Server, the XenData file system supports the usage of folders and files that contain these Mac specific characters. For files and folders containing Mac specific characters written to the XenData archive share, Macintosh users see the name as it was created.

**Important Limitation:** Windows users see the same name with any Mac specific characters replaced using a unicode to ANSI conversion. The Mac specific characters will not be displayed properly by Windows Explorer, Volume View, History Explorer and by the XenData Report Generator.

## Chapter 5. File Explorer and Windows Explorer Extensions

### 5.1 About File Explorer and Windows Explorer Extensions

XenData6 Server software extends the capabilities of File Explorer on Windows Server 2012 R2 and Windows Explorer on Windows Server 2008 R2 and Windows 7 to provide the following functionality.


- Smart copy and paste for optimized file restores from tape
- Flushing of files and folders
- Pre-fetching of files and folders
- History Explorer
- Volume View
- Enhanced properties

### 5.2 Smart Copy and Paste

The standard copy and paste operations available within File Explorer and Windows Explorer restore files from a tape archive in an order which does not take into account the location of the files on tape. When many files are being restored, this can cause considerable delays due to excessive tape seek and tape cartridge swap operations. The XenData Smart Copy and Paste operations are two alternative methods for restoring selected files from the archive in an optimized order which minimizes total restore time.


#### To Restore Files using Smart Paste

1. Open File Explorer or Windows Explorer
2. Select and then right-click on the required files and folders
3. Select **Copy**
4. Select the location to paste the copied files and folders
5. Right-click and select **XenData Smart Paste**



### To Restore Files using Smart Copy

1. Open File Explorer or Windows Explorer
2. Select, right-click and drag the selected files and folders to the required restore location
3. Unclick and then select **XenData Smart Copy**




### 5.3 Flushing

Selected files and the contents of selected folders can be flushed from disk using the Explorer Flush option. Flushing will only occur for files that have been successfully written to all designated replica tapes. The Explorer Flush option overrides the disk retention rules defined in the XenData Management Console. It does not cause flushing of the hidden resource fork files created by Apple Mac OS X clients.

Note that with all flushing operations, the file remains in the Windows file system; the flush operation causes the file to be removed from disk, but it is still visible and accessible to normal applications from tape. The Windows offline attribute is set for all flushed files.

#### To Flush Files

1. Open File Explorer or Windows Explorer
2. Select and then right-click on the required files and folders
3. Select **Flush**



File Explorer and Windows Explorer may spontaneously read one of the flushed files after a flush operation. If the disk retention rules, defined in the XenData Management Console, do not flush immediately after a file is closed, this will result in this file being pre-fetched.

#### 5.4 Pre-fetching

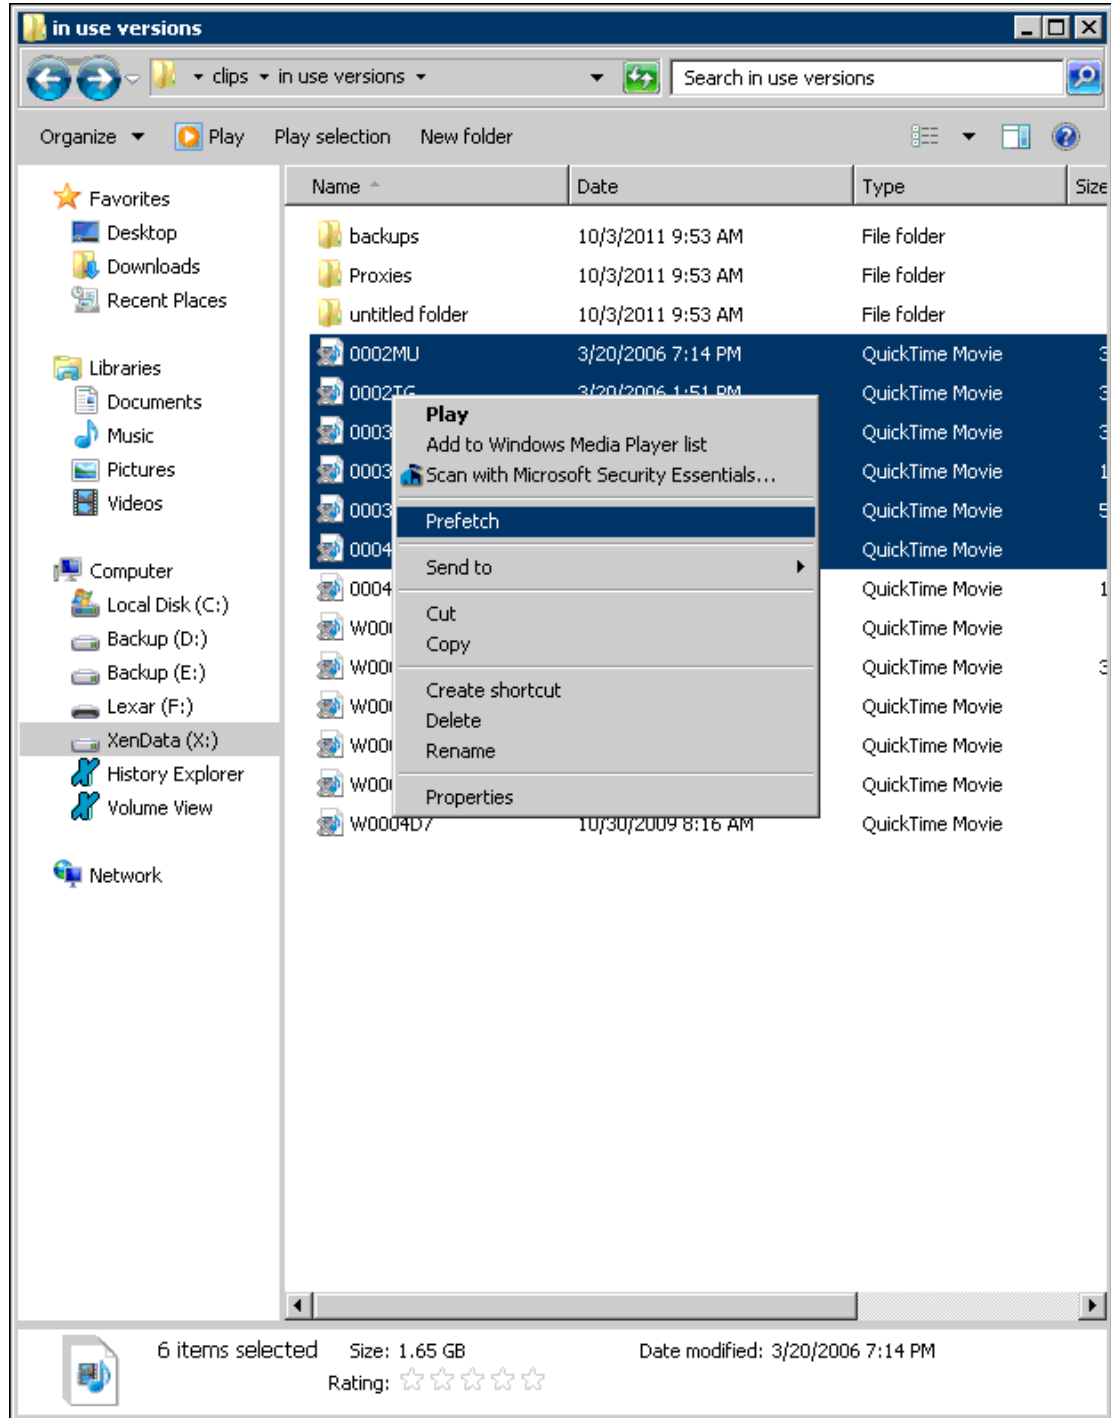
Selected files and the contents of selected folders can be pre-fetched from tape to disk cache using the Explorer Prefetch option.

The Explorer Pre-fetch option overrides the disk retention rules defined in the XenData Management Console. Pre-fetched files will not be automatically flushed unless read (in which case the disk retention rules defined for the applicable File Group apply) or manually flushed using Windows Explorer.

##### To Prefetch Files

1. Open File Explorer or Windows Explorer
2. Select and then right-click on the required files and folders
3. Select **Prefetch**





File Explorer and Windows Explorer may spontaneously read one of the pre-fetched files after a pre-fetch operation. If the disk retention rules, defined in the XenData Management Console, are set to flush after a file is closed, this will result in this file being flushed.

Note that if an individual file is selected, a recall option is also available. This also pre-fetches the selected file from tape to disk cache but additionally provides an on-screen display of any applicable error messages.

## 5.5 Volume View

Volume View is used to browse the contents of any tape Volume that the system knows about.

### To Browse with Volume View using Windows Explorer

1. Open File Explorer or Windows Explorer
2. Select Volume View in the left navigational pane
3. Browse the Volume View


## 5.6 History Explorer

History Explorer is used to obtain the complete history and status of any file that the system knows about. History Explorer lists all available versions of all files, all file instances and their media locations, including deleted and renamed files. It allows retrieval of all file versions and instances from disk or tape.

### To Browse with History Explorer

1. Open File Explorer or Windows Explorer
2. Select History Explorer in the left navigational pane
3. Browse the archive file system

Deleted files are shown with a grayed icon:




File properties may be launched by right-clicking on a file. The XenData tab will identify all file versions and by clicking a version to highlight it and then clicking open or copy, that file version may be opened or copied to another storage location.

### To View the Versions of a File

1. Right click on the file
2. Select Properties
3. Select the XenData tab

All versions of the file will be shown. An example with three file versions is illustrated below.

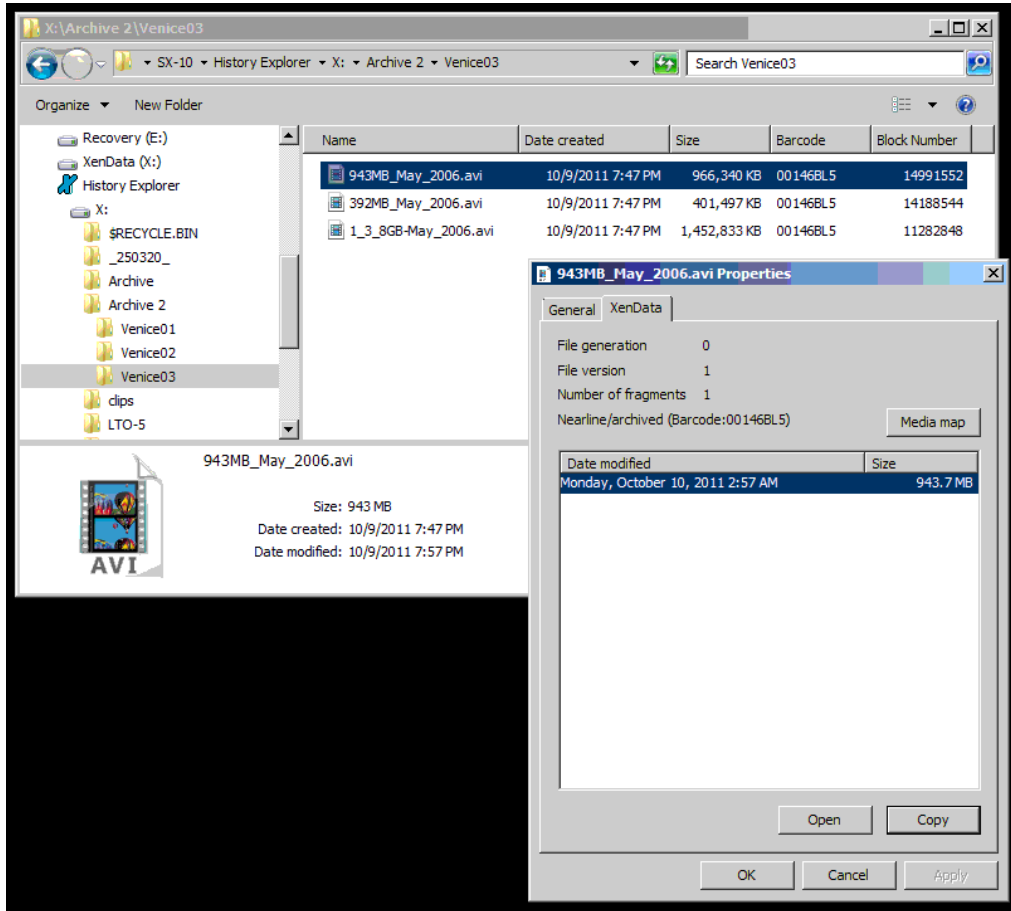


**To Open an Old Version of a File or a Deleted File**

1. View the file versions as described above.
2. Click on the required version of the file
3. Click **Open**

**To Restore an Old Version of a File or a Deleted File**

1. View the file versions as described above.
2. Click on the required version of the file
3. Click **Copy**
4. Use Windows Explorer to paste the file to the required storage location





## 5.7 Enhanced Properties via XenData Tab

Enhanced properties are available for the logical drive managed by XenData6 Server software as described below.

### To Obtain Enhanced Properties

1. Open File Explorer or Windows Explorer
2. Right click on the logical drive letter under XenData control
3. Select **Properties** and then select the **XenData** tab



The following properties are displayed.

- **Software version** The installed version of XenData6 Server software.
- **Total managed slots/cartridges** The total number of slots in attached tape libraries and tape drives plus the number of managed offline cartridges.
- **Full cartridges** The total number of full cartridges including finalized cartridges that are only partially filled and offline cartridges.
- **Writable cartridges** The number of writable cartridges, excluding blank cartridges and offline cartridges.
- **Unwritable cartridges** The number of non-writable cartridges including offline cartridges.
- **Blank cartridges** The number of blank cartridges in attached tape libraries and tape drives.
- **Empty slots** The number of empty slots in attached tape libraries and tape drives.
- **Offline cartridges** The number of offline tape cartridges managed by the system.
- **Cache volume capacity** The capacity of the magnetic disk volume under XenData control.
- **Cache volume free space** The free space available on the magnetic disk volume under XenData control.

## Chapter 6. Metadata Backup

### 6.1 Overview

The XenData Metadata Backup program deals with:

- **File system metadata** which is stored on the disk logical drive dedicated for use by the XenData6 Server software.
- The **XenData State File** which contains cartridge information and the XenData Management Console settings, including File Group and Volume Set configuration settings.

If the system has to be rebuilt, perhaps due to a RAID failure, the file system metadata may be rebuilt on disk without using the XenData Metadata Backup program by using the Build Catalog, Import Folder Structure and Import Data functions available from the XenData Management Console. However, this can be a lengthy process for a system with a large number of Volumes. Alternatively, by using the XenData Metadata Backup program, the system can be rebuilt more rapidly.

The XenData Metadata Backup program speeds up the process of rebuilding the magnetic disk logical drive by restoring the file system metadata and XenData state file to the condition they were in at the time of the backup. This means that the Build Catalog and Import Folder Structure functions need only be used for Volumes which have been written to since the latest backup.


### 6.2 Using the Program

#### 6.2.1 Starting the Program

For Windows Server 2008 R2 and Windows 7:

1. Click **Start**
2. Click **Programs** (or **All Programs**)
3. Click **XenData**
4. Click **XenData Metadata Backup**

For Windows Server 2012 R2:


1. Click on the **Start** button, located in the lower-left hand corner of the screen
2. Click the arrow  near the lower-left corner of the Start screen.
3. Click on **XenData Metadata Backup** shown under the **XenData** category

#### 6.2.2 Selecting Backup or Restore

The XenData Metadata Backup program performs two types of operation:

- **Make backup** - Make a backup of the metadata in the system in its current state. See '**Making a Pre-defined backup**' or '**Making a Custom Backup**' below.
- **Restore from backup** - Restore information from a backup file into the current XenData logical drive. See '**Restore from backup**' below.


Select the desired option and click **Next** to continue.



### 6.2.3 Making a Pre-defined Backup

The instructions in this section describe how to perform a backup using a pre-defined backup type. The section **Making a Custom Backup** shows how to exert more control over what is backed up (for example to exclude a temporary directory).

Having started the XenData Metadata Backup program and selected **Make backup**, click **Next**.




There are two pre-defined backup types. **All metadata** will back up all of the file system metadata, and **All metadata and XenData state file** will also include the XenData state file.

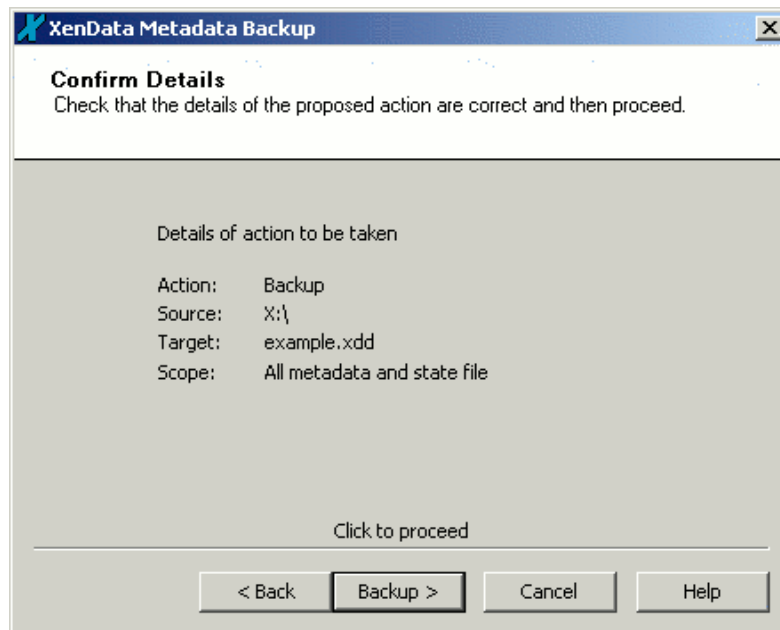
1. Select **All metadata** or **All metadata and XenData state file** as appropriate.
2. Click **Next** to continue.



(Note that the option **Let me choose what to back up** provides more control over what is backed up than the two pre-defined backup types. For example, a folder only used for temporary files may be excluded from the backup as the files it contains will not be necessary again. See **Making a Custom Backup** for more information.)




1. Verify the XenData logical drive to be backed up. Although XenData6 Server only supports one XenData logical drive per system, you should verify that it has been correctly detected as the source for the backup.
2. Specify the output path and file name. The output file name should be inserted in the **Backup to** edit box. Click **Browse** to assist in specifying the path and filename.
3. Click **Next** to continue.



The next page presents the details of the backup, and gives the option to go back and correct if necessary.

1. Verify the backup details.
2. Click **Backup** to perform the backup.

A progress dialog appears that shows the backup progress, as illustrated below.



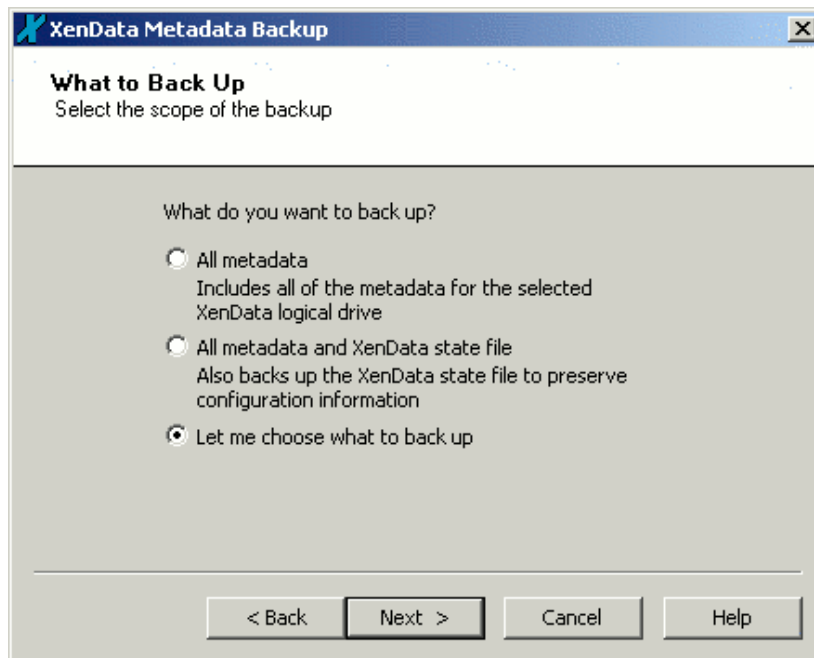
If the backup completed successfully, you will be presented with a confirmation page saying **Backup Complete**. Click **Finish** to dismiss the dialog and exit the program.

#### 6.2.4 Making a Custom Backup

The instructions in this section describe how to perform a backup with control over what is included in the backup.

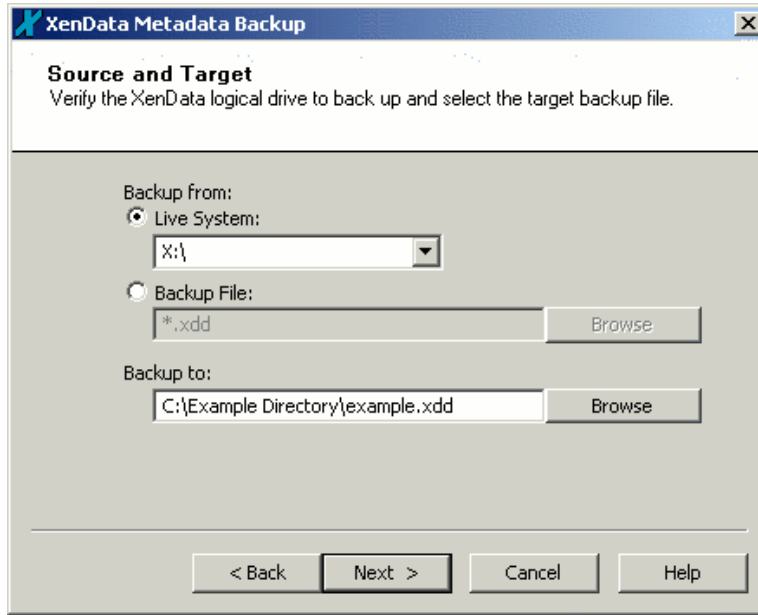
For example, a folder only used for temporary files may be excluded from the backup as the files it contains will not be necessary again. It also allows the creation of a sub-backup. This refers to creating a new backup file from an existing backup where the new backup contains only selected folders from the existing backup file.

Open XenData Metadata Backup. Select **Make backup** and click **Next** on the starting page as described in **Backup and Restore**.



The option **Let me choose what to back up** provides control over what file system metadata gets backed up, and whether the XenData state file is included.


1. Select **Let me choose what to back up**.
2. Click **Next** to continue.



If a sub-backup of an existing backup file is being made, an already existing backup file should be selected as the source. This cannot then be used as the target backup file. The **Browse** buttons can be used to assist in specifying the filenames.

Although XenData6 Server only supports one XenData logical drive per system, you should verify that it has been correctly detected as the source for the backup.


1. Select **Live System** or **Backup File** as appropriate.
2. Either verify the XenData logical drive to backup from or specify the backup file to use as a source as appropriate
3. Specify the output file name.
4. Click **Next** to continue.



A folder which is to be included in the backup is marked with a black check mark, and one which is to be ignored is left unchecked. A directory whose presence will be recorded but for which no file system metadata will be saved is marked with a 'grayed out' check mark.

Clicking on the "+" sign expands a subdirectory tree, and clicking on a "-" sign contracts it. Selecting a folder in the tree, to include all of the files in it, causes all of its parent folders to be selected as required to enable the selection to take place.


1. Select and deselect folders in the tree as appropriate to indicate what should be backed up.
2. Click **Finish** to continue.




This page presents the details of the backup, and gives the option to go back and correct if necessary.

1. Verify the backup details.
2. Click **Backup** to perform the backup.

A progress dialog appears that shows the backup progress, as illustrated below.



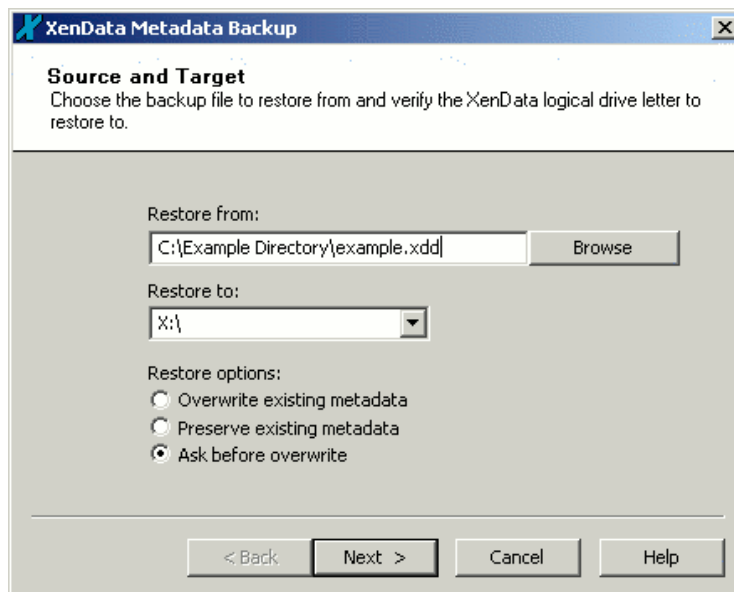
If the backup completed successfully, you will be presented with a confirmation page saying **Backup Complete**. Click **Finish** to dismiss the dialog and exit the program.



### 6.2.5 Restore from Backup

The instructions in this section describe how to restore a selection of the file system metadata in a backup file onto a live system, and/or restoring the XenData state file.

Either start the XenData Metadata Backup program, select **Restore from backup** and click **Next** on the starting page as described in **Backup and Restore**, or double click on a XenData backup file (\*.xdd) to enter the **Restore from backup** section.



If the restore section has been accessed from the start page, the **Back** button will be enabled.

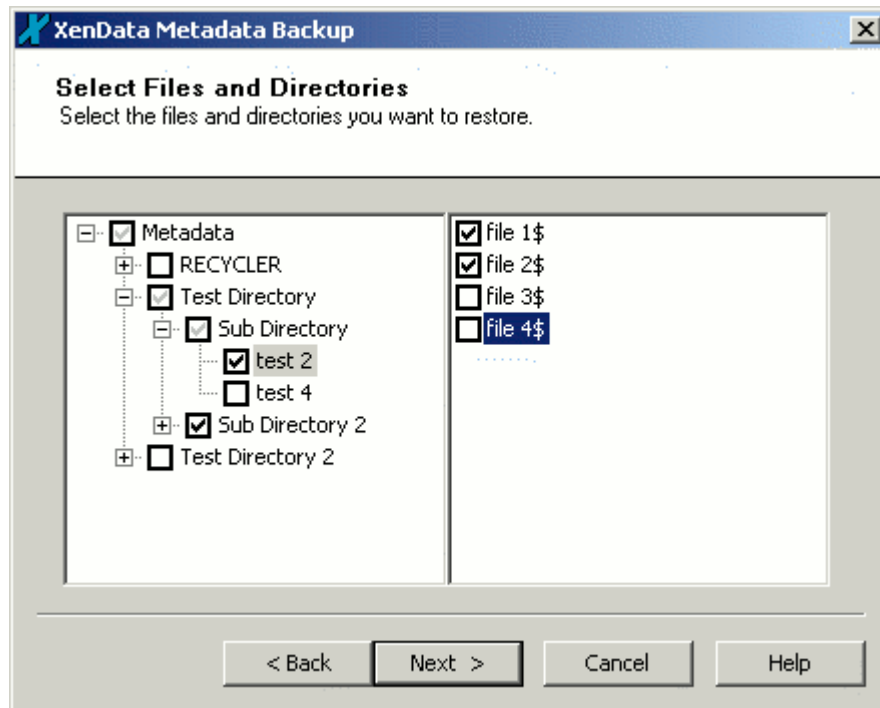
There are three restore options

- **Overwrite existing metadata** always writes metadata from the file into the XenData logical drive, ignoring metadata already present.
- **Preserve existing metadata** will only write metadata for a particular file from the backup if there is none already present on the XenData logical drive for that file.

- **Ask before overwrite** will ask whether to overwrite existing metadata for each file whose metadata already exists, providing options to overwrite all of a certain type (for example, the metadata on the XenData logical drive is currently invalid).

Although XenData6 Server only supports one XenData logical drive per system, you should verify that it has been correctly detected as the target for the restore. Use the file dialog opened by clicking **Browse** to assist in specifying the filename of the source backup file.


1. Specify the input backup file to restore from, or verify that the correct filename has been determined automatically.
2. Verify the XenData logical drive to restore to.
3. Select the desired **restore option**.
4. Click **Next** to continue.



A folder or file which is to be restored is marked with a black check mark, and one which is to be ignored is left unchecked. A directory which needs to be traversed to reach checked items, but which will not itself be included is marked with a 'grayed out' check mark. If a directory is selected, the files within it are all selected by default, unless manually deselected.

Clicking on the "+" sign expands a subdirectory tree, and clicking on a "-" sign contracts it. Selecting a directory or file in the tree, in order to include it in the restore, causes all of its parent directories to be selected as required to enable the selection to take place.



1. Select and deselect directories and files in the tree as appropriate to indicate what should be restored.
2. Click **Next** to continue.



This page presents the details of the restore, and gives the option to go back and correct if necessary.

1. Verify the restore details.
2. Click **Restore** to perform the restore.

A progress dialog will appear so that one can check that the restore is progressing:

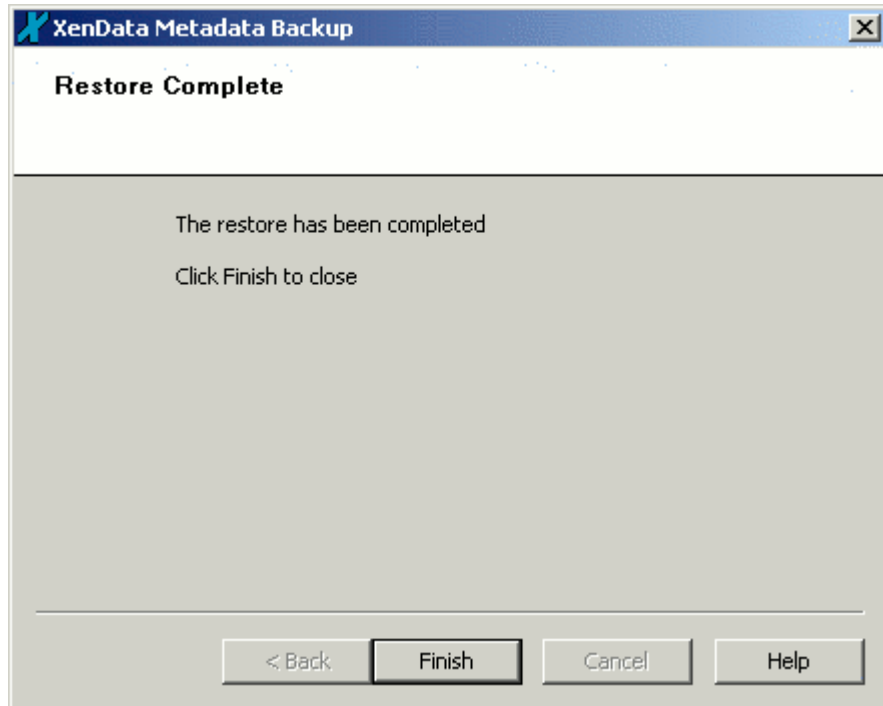


If **Ask before overwrite** was specified in the first step, dialog boxes similar to the one shown above may appear, asking if existing metadata should be overwritten, and giving a category of file to consider - in this case where the original metadata is inconsistent. This gives the option to deal with these cases on a file by file basis (yes/no) or to specify what action should be taken for all files of this type (Yes for All/No for All), which prevent further dialogs appearing.

1. Click **Yes** or **No** to choose whether to overwrite the file system metadata for the current file.
2. Click **Yes for All** or **No for All** to choose whether to overwrite the file system metadata for all files in the same category.

**Note:** If the metadata on disk for a file is exactly the same as that in the backup file, no overwrite dialog will be displayed, no change is necessary and the file will be silently skipped.

If the restore completed successfully, you will be presented with a confirmation page saying **Restore Complete**. Click **Finish** to dismiss the dialog and exit the program.





## Chapter 7. Scheduler

### 7.1 Overview

The XenData Scheduler can be used to schedule the following task types:

- **Metadata Backup** This allows scheduling of full metadata backups including backup of the XenData state file. It does not support scheduling of custom backups. Metadata backup and restore is described in Chapter 6.
- **Deferred Writing to Data Tape** This defers writing of files to the primary tape cartridges of one or more Volume Sets and then enables writing for a scheduled time period. It is useful for making tape drives available only for restoring files during times of frequent restores. Note the scheduler does not defer writing to replica tape cartridges; this must be configured as described in Section 3.2.7.


### 7.2 Using the Program

#### 7.2.1 Starting the Program

For Windows Server 2008 R2 and Windows 7:

1. Click **Start**
2. Click **Programs** (or **All Programs**)
3. Click **XenData**
4. Click **XenData Scheduler**


For Windows Server 2012 R2:

1. Click on the **Start** button, located in the lower-left hand corner of the screen
2. Click the arrow  near the lower-left corner of the Start screen.
3. Click on **XenData Scheduler** shown under the **XenData** category

This launches the Scheduler status display.

#### 7.2.2 To Add a Task

Start the Scheduler program and click on 'Add New Task' and then select the type of task from the drop down menu as shown below.



### 7.2.3 To Edit a Task

1. Start the XenData Scheduler
2. Select a Task from the list with Status 'Idle'
3. Click the **Edit** button

### 7.2.4 To Delete a Task

1. Start the XenData Scheduler
2. Select a Task from the list with Status 'Idle'
3. Click the **Delete** button

### 7.2.5 To Immediately Run a Task

1. Start the XenData Scheduler
2. Select a Task from the list with Status 'Idle'
3. Click the **Run Now** button

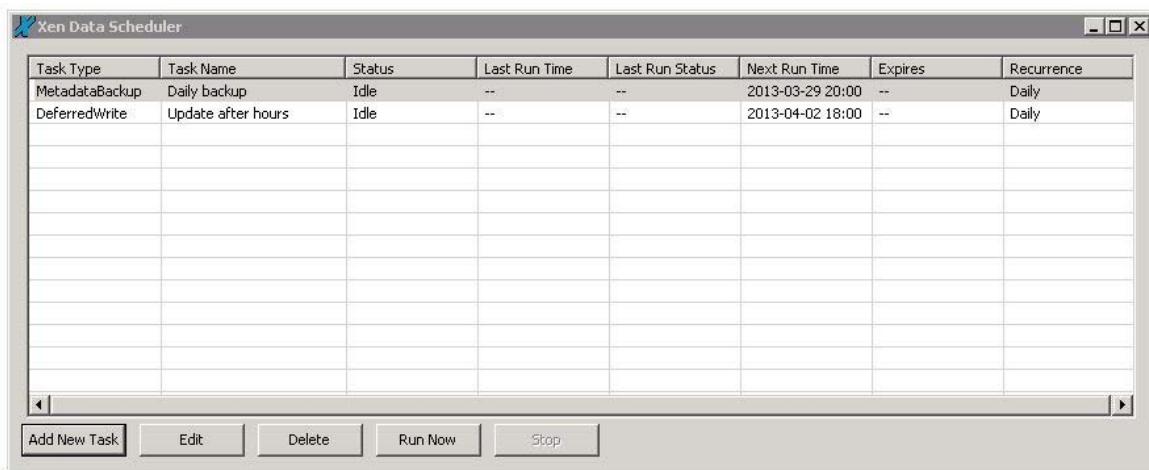
### 7.2.6 To Immediately Stop a Task

1. Start the XenData Scheduler
2. Select a Task from the list with Status 'Running'
3. Click the **Stop** button

**NOTE** When a Metadata Backup Task is stopped by force (i.e. by using the **Stop** button) its 'Last Run Status' is set to 'FAIL' and no backup file is created.

## 7.3 About the Scheduler Status Display

An example of the Scheduler status display is shown below.



The display columns are described below.

**Task Type** shows the task type.

**Task Name** shows the name given to this task. It is an optional parameter and it can be empty.

**Status** shows the current task status. The status can be:

- **Idle** – The task is not running. In this state one can 'Edit', 'Run Now' or 'Delete' the task
- **Running** – The task is running. In this state one can only 'Stop' the task
- **Locked** – The task is being edited by another user. The task remains locked until the editing is completed.

**Last Run Time** This shows the most recent date and time when the task was run. '--' indicates that the task was never run.

**Last Run Status** This shows the result of the last task run. The status can be:

- '--' – The task has not yet been run.
- **OK** – The task ran and finished successfully.
- **FAIL** – The task failed.
- **Paused OK** – The task was stopped before it finished.

**Next Run Time** This shows the date and time when the task will be run again. '--' indicates that the task will not be run again.

**Expires** This shows the date and time when the task recurrence expires. After this date the task will not be run again. '--' indicates that the task never expires.

**Recurrence** This shows the task recurrence setting. Recurrence can be:

- **None** - Task is only run once.
- **Daily** - Task is run once per day until it expires
- **Weekly** - Task is run once per week until it expires
- **Monthly** - Task is run once per month until it expires.

## 7.4 How to Configure a Metadata Backup Task

An example of a Metadata Backup Task configuration is shown below.

**Recurrence** This shows the task recurrence setting. Recurrence can be:

- **None** - Task is only run once.
- **Daily** - Task is run once per day until it expires
- **Weekly** - Task is run once per week until it expires
- **Monthly** - Task is run once per month until it expires.

**Expire** This shows the date and time when the task recurrence expires. After this date, the task will not be run again. '--' indicates that the task never expires.

**Start** This is used to set the date and time for the first run of the task.

**Task Name** This is used to set the task name. This is an optional parameter. It can be blank.

**Chose directory path for backup** This is used to define the folder location for the backup files. Every time the task is run successfully, a backup file is created with name 'YYYYMMDDHM.xdd'. For example 'f:\metabackup\201207251905.xdd'

**Delete previous backups** This option deletes all previous metadata backup files (i.e. all files that end with .xdd) in the backup file folder location.

**Save** This saves the current configuration and returns to the scheduler status display.

**Cancel** This ignores all the changes and returns to the scheduler status display.

## 7.5 How to Configure a Deferred Write Task

An example of a Deferred Write Task configuration is shown below.

**Recurrence** This shows the task recurrence setting. Recurrence can be:

- **None** - Task is only run once.
- **Daily** - Task is run once per day until it expires
- **Weekly** - Task is run once per week until it expires
- **Monthly** - Task is run once per month until it expires.

**Expire** This shows the date and time when the task recurrence expires. After this date, the task will not be run again. '--' indicates that the task never expires.

**Start** This is used to set the date and time for the first run of the task.

**Task Name** This is used to set the task name. This is an optional parameter. It can be blank.

**Force stop task if it runs longer than** This is used to set the maximum run time for the task.

**Volume Sets with Deferred Write Enabled** The list shows the Volume Sets for which deferred writing to the primary tape has been enabled. Only the selected Volume Sets in the list are controlled by this task (i.e. only the list items with a check as shown above). Note that at least one Volume Set must be selected. To disable deferred writing for a Volume Set, click on the Volume Set in the list and then click on the '--->' button.

**Volume Sets with Deferred Write Disabled** The list shows the Volume Sets for which deferred writing has not been enabled. To enable deferred writing, click on the Volume Set in the list and then click on the '<---' button.

**Save** This saves the current configuration and returns to the scheduler status display.

**Cancel** This ignores all the changes and returns to the scheduler status display.

## Chapter 8. System Recovery

### 8.1 System Rebuild from Tape

#### 8.1.1 About File System Rebuild From Tape


The system provides a facility to entirely rebuild itself from tape, for example in the event of a disk failure. This capability may be used after a XenData metadata restore using the XenData Metadata Backup utility described in Chapter 6.

If an imported tape cartridge belongs to a Volume Set that has not previously been defined on this particular server, the new Volume Set will be created and the tape cartridge will be allocated to it automatically. Additionally, if the tape cartridge was written using LTFS or was written using TAR and has been finalized, the tape catalog will be automatically copied to the disk. However, the files present on the cartridge will not be visible in the file system and available for restore until the **Import Folder Structure** or **Import Data** operation is performed.

Selecting **Import Folder Structure** loads file and folder information (file system metadata) into the system, making the entire file system visible to users, but it does not restore the actual file data to disk, so any files that are subsequently read from the system will be restored from tape at the time they are read. The Import Data option also loads file and folder metadata but additionally, it loads file data into the system, in accordance with the File Group rules and the Disk Retention Rules for written files.

#### 8.1.2 To Rebuild a File System from Tape

1. Open the XenData Management Console.
2. Import the tape cartridges into the library or insert into a stand-alone tape drive. The system automatically creates one or more Volume Sets containing the newly imported cartridges. Additionally, if the cartridges were formatted with LTFS or are finalized TAR formatted cartridges, their tape catalogs will be automatically copied to the disk.
3. At this step in the process, you may choose to use **Import Folder Structure** which will make all the files on finalized tapes visible in the file system and available for restore. The Import Folder Structure operation may be performed for individual Volumes, for Volume Sets or for all Volume Sets by right clicking on the appropriate item in the left pane of the XenData management Console and selecting **Import Folder Structure**.
4. If one or more of the tape cartridges was written using TAR and was not finalized, select 'Volume set configuration' in the XenData Management Console. Right click and select **Build Missing Catalogs**. (Note that this causes scanning of the non-finalized tapes and may take several hours.)



Following successful creation of all the Tape Catalogs, right click on 'Volume set configuration' and select **Import Folder Structure**. This makes all the files visible in the logical drive under XenData control and available for restore.

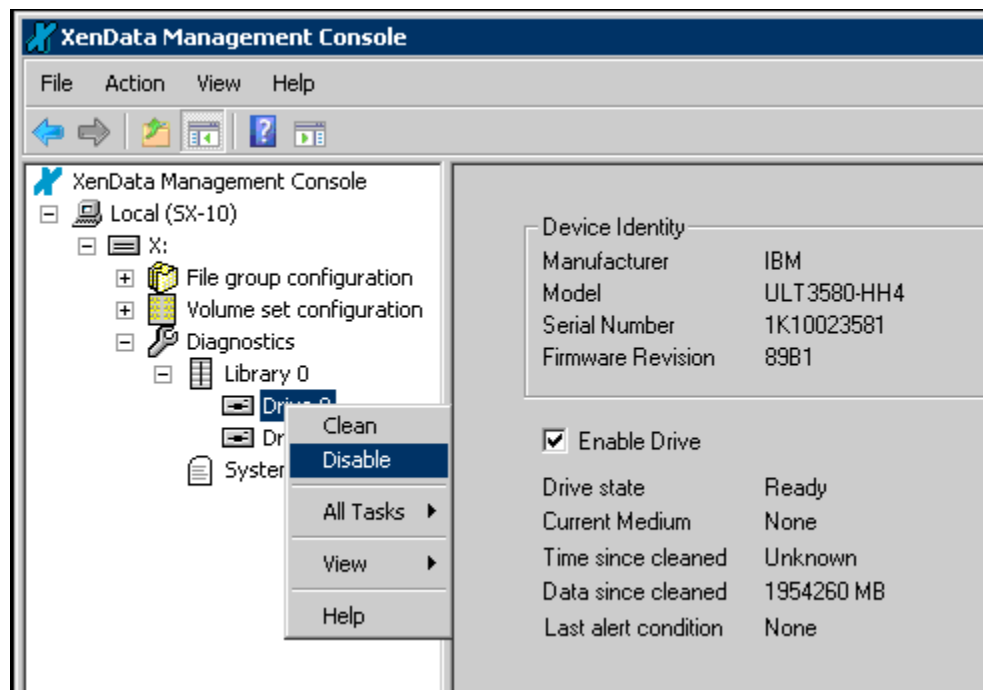
To cache files on disk, use the **Import Data** option which creates instances of files on the disk under XenData control in accordance with the File Group rules and the Disk Retention Rules for written files. (Note that this causes scanning of the tapes and may take several hours.)

## 8.2 In Case of Hardware Failure

Occasionally, it may be necessary to temporarily disable one or more tape drives or libraries, perhaps for preventive maintenance. XenData6 Server software allows the system administrator to selectively disable hardware while the remainder of the system continues running.

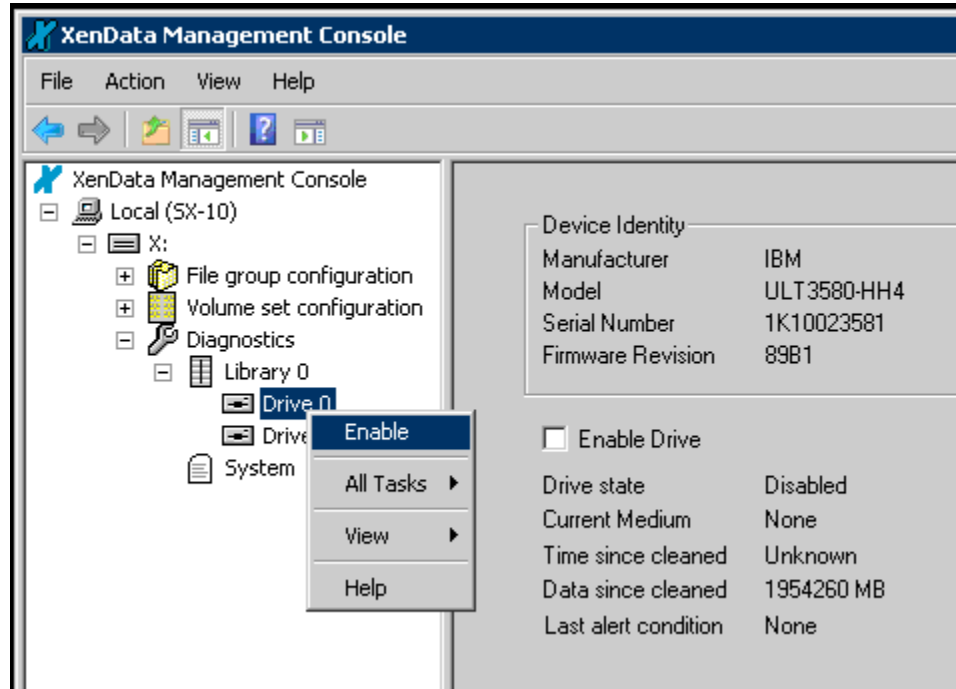
### 8.2.1 To Disable a Tape Drive or Library

1. Open the XenData Management Console.
2. Navigate to the **Diagnostics** section
3. Right-click on the component and select '*Disable*'.



### 8.2.2 To Re-enable a Tape Drive or Library

1. Open the XenData Management Console.
2. Navigate to the **Diagnostics** section
3. Right-click on the component and select '*Enable*'



### 8.2.3 Options In Case of Tape Library or Stand-Alone Drive Failure

In normal operation, XenData6 Server software writes files to tape immediately after they have been written to disk cache.

When all cartridges in a Volume Set become full or unavailable, the system response to an attempt to write additional files to the Volume Set depends on the Volume Set configuration settings. All tapes in a Volume Set may become full or unavailable for a number of reasons:

- use of stand-alone tape drives and the inserted tape or tapes have simply become full;
- use of a tape library with insufficient blank tapes;
- a failure of the tape drive or library hardware; or
- a media error.

If the 'Write to disk if no writable media are available' option has been enabled and all tapes in a Volume Set become full or unavailable, the system automatically enters the Pending Write Mode and will accept more data. However, if the 'Write to disk if no writable media are available' option has not been enabled, the system will not accept any more data and will report "disk full" when an attempt is made to write to the Volume Set.

When the system enters the 'Pending Write Mode', it defers writing to tape and continues writing to disk. When free space within the Volume Set becomes available again, the system automatically 'catches up' and writes the deferred files to tape. When the system is in the Pending Write Mode, a comprehensive set of warning messages are sent to the XenData Archive Event Log. These include notification of entering and leaving the Pending Write Mode and warnings as disk space becomes exhausted. The XenData Alert Module may be configured to provide notification via email and / or on-screen message of these warning messages.



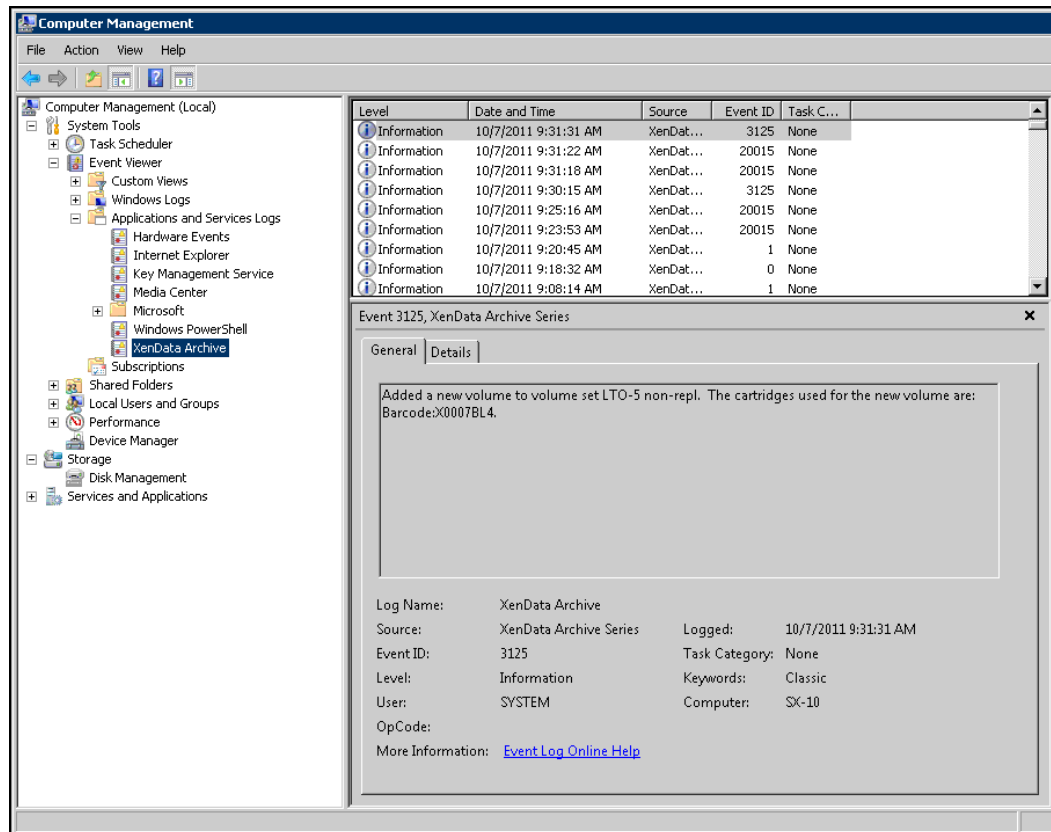
## Chapter 9. Diagnostics & Maintenance

### 9.1 XenData Event Log

Whenever XenData6 Server software encounters an unexpected error condition, it puts a message in the Windows Event Log and generates a Trace Log file. In general, if the system is not behaving as expected, the Event Log is the first place that a system administrator should look.

#### To Open the XenData Event Log

1. Right-click on "Computer"
2. Select **Manage**
3. Navigate to the XenData section of the Event Viewer in the Computer Management console, as shown below.




### 9.2 XenData System Trace Log

It is sometimes useful to be able to see what is happening internally within the system. The System Trace Log allows the system administrator to examine a trace of all actions performed by the system.

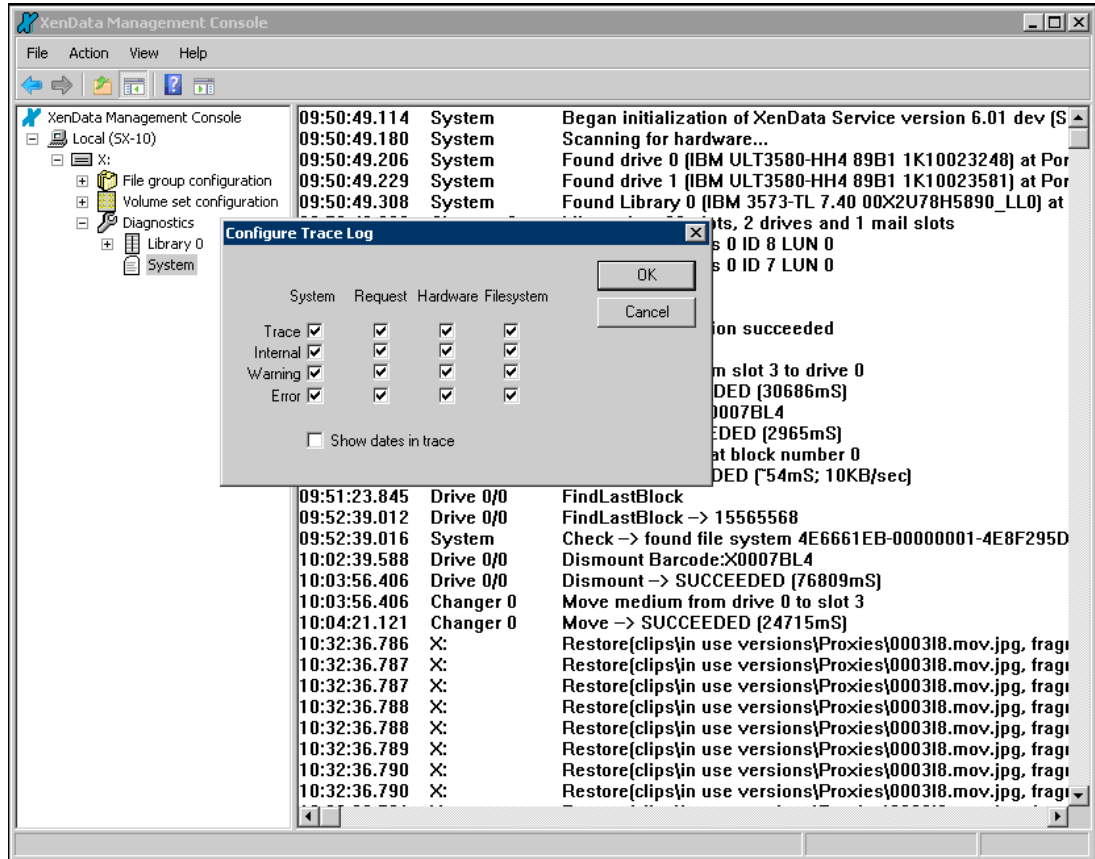
#### 9.2.1 To Open the Trace Log

1. Open the XenData Management Console.
2. Navigate to the **Diagnostics** section.
3. Click on the System icon to open the trace log in the right pane of the window.



### 9.2.2 To Change the Level of Detail in the Trace Log

The trace log can generate a large amount of information, which can scroll past very quickly. Right-clicking on the System icon reveals the "Configure" option, which brings up a dialog box that allows configuration of the components to be traced and the level of detail of the trace:



### 9.2.3 Automatic Generation of Trace Files

Whenever XenData6 Server software encounters an unexpected error condition, it puts a message in the XenData Event Log and may generate a trace file. The trace file contains a record of what the system was doing in the time preceding the error, and is especially useful to assist XenData support personnel in determining the cause of a problem.

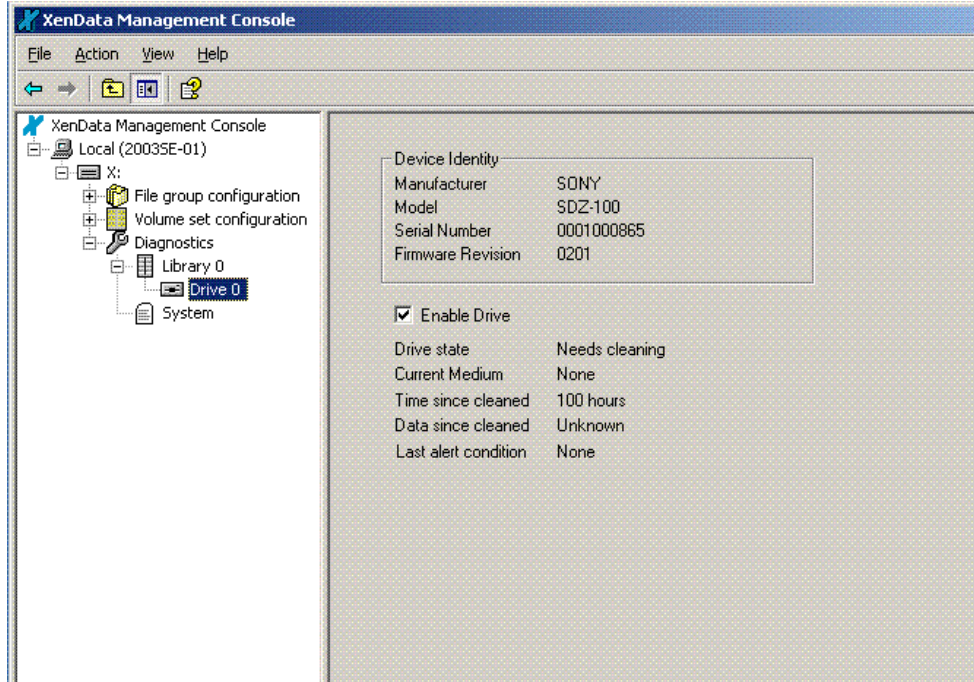
Trace files have the extension .xdt and are stored in the \XenDataLog folder of the system boot drive. They are saved in a compressed format to make them easier to transmit by email. A supplied utility (XDTraceViewer.exe) is required to open and read the contents of a trace file.

## 9.3 To View Tape Drive and Library Diagnostic Information

XenData6 Server software provides diagnostic information about the hardware under its control, and allows the system administrator to selectively disable hardware that may need attention. This information can be found in the Diagnostics section of the XenData Management Console.

### To Open the Hardware Diagnostics

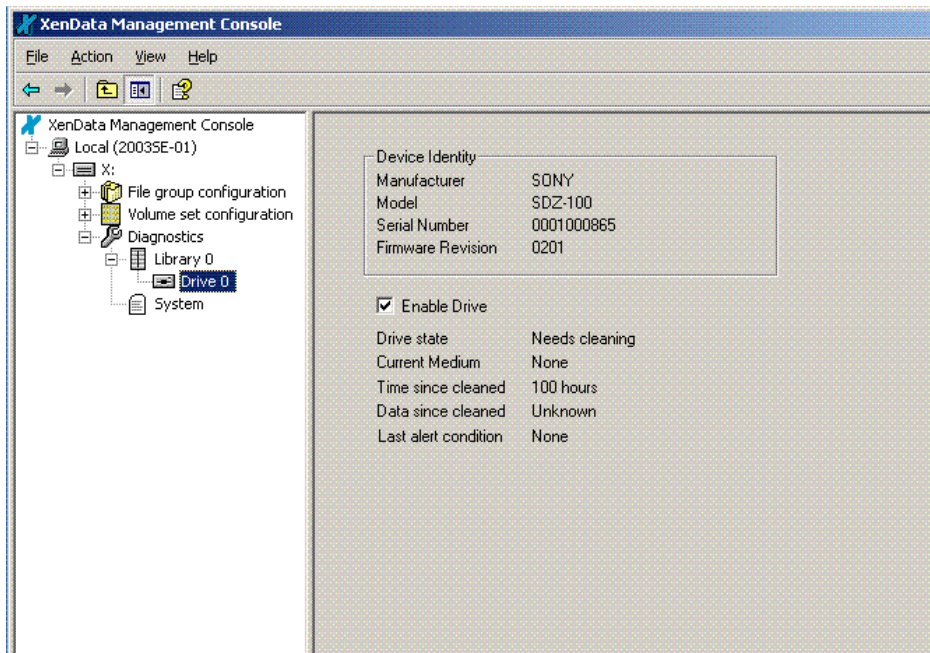
1. Open the XenData Management Console.
2. Navigate to the **Diagnostics** section.
3. Click on a hardware component to display information about that component.



## 9.4 Tape Drive Cleaning

### 9.4.1 About Tape Drive Cleaning


Most types of tape drive require periodic cleaning in order to maintain peak performance. Normally, a drive will automatically issue a request for cleaning at the appropriate time and in the case of drives within tape libraries, XenData6 Server software responds to these requests by inserting a cleaning cartridge in the drive. In the case of stand-alone tape drives or if no cleaning cartridge is available in a tape library, the system will put a message in the Windows Event Log and in the XenData Management Console identifying that the drive requires cleaning.



If required, the system administrator can manually clean a drive within a tape library as follows:

#### 9.4.2 To Manually Clean a Drive within a Library

1. Ensure that a cleaning cartridge has been imported into the library.
2. Open the XenData Management Console.
3. Navigate to the **Diagnostics** section.
4. Expand the Library icon and find the appropriate drive
5. Right-click on the drive and select 'Clean'




## 9.5 Tape Cartridge Alert State

### 9.5.1 About the Tape Cartridge Alert State

Certain error conditions, such as media errors, leave the affected tape (or, in the case of replicated tapes, the affected Volume) in an "Alert" state. The system does this to protect data by ensuring that it will not attempt to write to tape cartridges that have problems associated with them. When such an error occurs, the system also puts a message in the event log. After consideration of the error message in the event log, the system administrator may decide that it is appropriate to ignore the error and continue to use the affected tape(s). The system provides a mechanism to reset the alert state, as follows:

#### 9.5.2 To Clear the Alert State for a Tape

1. Open the XenData Management Console
2. Select the tape in the XenData Management Console
3. Right click on the tape and select **Reset alert**. Note that the **Reset Alert** option is only available for tapes that are in the "Alert" state.



In some cases, it may not be possible to reset the alert state because of the severity of the original error condition. In these cases, files written to the tape cartridge prior to the error condition will still be readable.

However, additional files cannot be written to the Volume and a new Volume must be added to the Volume Set to allow writing of additional files.

## Chapter 10. Reports

### 10.1 About the Report Generator

The report generator provides reports for an archive running XenData6 Server software. Available reports are:

- **File Search** This report lists files that match a search term and identifies the tape cartridges on which they are stored.
- **Recoverable Space** This report lists the amount of space that can be recovered from a Volume using a Repack operation.
- **Tape Cartridge Contents** This displays the files stored on a selected tape cartridge.
- **Unarchived Files** This report displays files that are not written to tape as defined by the current File Group rules.
- **Volume Contents** This displays the files stored on a selected Volume.

Reports may be saved in a number of different formats, including as a text file that may be imported into Microsoft Excel.


### 10.2 Starting the Report Generator

To start the XenData Report Generator

For Windows Server 2008 R2 and Windows 7:

1. Click **Start**
2. Click **Programs**
3. Click **XenData**
4. Click **XenData Report Generator**


For Windows Server 2012 R2:

1. Click on the **Start** button, located in the lower-left hand corner of the screen
2. Click the arrow  near the lower-left corner of the Start screen.
3. Click on **XenData Report Generator** shown under the **XenData** category


### 10.3 Creating, Saving and Restoring Reports

#### 10.3.1 To Create a Report

Start the Report Generator program and from the initial page, select **File** and then **New** as shown below.



Then select the required report type from the drop down menu as shown below.




Please refer to the applicable section below for instructions on the selected report type.

### 10.3.2 To Save a Report

A report can be saved in three different formats: Report Generator format (.XRG), tab delimited plain text (.txt) or XML. The XRG format is the only format which can be displayed by the Report Generator. The text format is useful for exporting the results to Microsoft Excel.

To save a report, select 'File' and then 'Save As' as shown below.



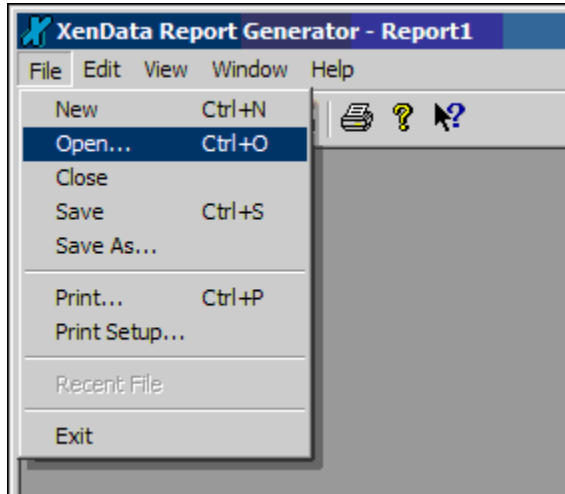
Then browse to the required location, select the file name and format and then click 'Save'.

### 10.3.3 To Display a Saved Report

The Report Generator will display reports saved in the XRG format only.

Start the Report Generator program and from the initial page, select 'File' and then 'Open' as shown below.






Then browse to the location of the saved report, select the required XRG file and open it.

## 10.4 File Search Reports

### 10.4.1 To Run a File Search Report



The File Search Report lists files archived to tape that match a search term and identifies the tape cartridges on which they are stored. The search may be limited to one selected Volume Set or may include all Volume Sets. The displayed report can be filtered to show one of the following:

- **All Files** This displays all files on tape including deleted files, old versions of files and renamed files.

- **Only Current Files** This displays only the current files. These are files that can be accessed via the archive file system interface and excludes deleted files, old versions of files and renamed files.
- **Only Deleted Files** This displays only deleted files.

You can search for a file name with specified text or by using a Regular Expression. When **File Name Text Search** is chosen, the search option supports wildcards.

Having selected the Volume Set, the filtering options and search type and term, click **Get Report**.

**Note:** A File Search Report will search only in the Volumes that have a tape contents catalog file cached on the system. Volumes created with versions of XenData Archive Series software prior to 4.50 do not automatically have tape catalogs. Also, tapes that have not been finalized and have been imported from another archive system will not automatically have tape catalogs.

### 10.4.2 Interpret a File Search Report

An example of a File Search Report is shown below.

**XenData Report Generator - [Report1]**

File Edit View Window Help

Report Type: FileSearch  
 Volume Set: All Volume Sets  
 Showing: All Files  
 Search Type: File Name Text  
 Search Term: venice

No	File Name	Barcode	Generation	Version	Status
1	/Venice/Venice1/943MB_May_2006.avi	Barcode:X0015A	0	1	Current
2	/Venice/Venice1/1_3_8GB-May_2006.avi	Barcode:X0015A	0	1	Current
3	/Venice/Venice1/392MB_May_2006.avi	Barcode:X0015A	0	1	Current
4	/Venice/Venice2/1_3_8GB-May_2006.avi	Barcode:X0015A	0	1	Current
5	/Venice/Venice2/392MB_May_2006.avi	Barcode:X0015A	0	1	Current
6	/Venice/Venice2/943MB_May_2006.avi	Barcode:X0015A	0	1	Current
7	/Venice/Venice3/943MB_May_2006.avi	Barcode:X0015A	0	1	Deleted
8	/Venice/Venice3/1_3_8GB-May_2006.avi	Barcode:X0015A	0	1	Deleted
9	/Venice/Venice3/392MB_May_2006.avi	Barcode:X0015A	0	1	Deleted
10	/Venice/Venice4/1_3_8GB-May_2006.avi	Barcode:X0015A	0	1	Overwritten
11	/Venice/Venice4/392MB_May_2006.avi	Barcode:X0015A	0	1	Overwritten
12	/Venice/Venice4/943MB_May_2006.avi	Barcode:X0015A	0	1	Overwritten
13	/Venice/Venice4/943MB_May_2006.avi	Barcode:X0015A	0	2	Current
14	/Venice/Venice4/1_3_8GB-May_2006.avi	Barcode:X0015A	0	2	Current
15	/Venice/Venice4/392MB_May_2006.avi	Barcode:X0015A	0	2	Current
16	/Venice/Venice10/1_3_8GB-May_2006.avi	Barcode:00147BL5	0	1	Current
17	/Venice/Venice10/392MB_May_2006.avi	Barcode:00147BL5	0	1	Current
18	/Venice/Venice10/943MB_May_2006.avi	Barcode:00147BL5	0	1	Current

The File Search Report lists files archived to tape that match the specified search criteria. The display columns are described below.

**No.** This is the number of the file in the display sorted by either date or file name, as defined by the **Sort** by selection.

**File Name** This is the file name including full path from the root of the archive logical drive letter.

**Generation** When a file of a given name and path is first created, it is termed generation 0. Every time the file is deleted or renamed and then a new file of the same name is created, the system increments the generation number. Note that each time the generation number is incremented, the version sequence starts again, with version 1 of the new file being the first that contains data.

**Version** If a file is updated with a newer version by overwriting or appending, the XenData software assigns a new version number. A file's version number increases by one every time it has data written to it. Note that the version number does not increase for every individual write operation, just for every file open that is followed by a write. Version 0 of a file never contains any data; the first time an application writes to the file, the version number is incremented to 1.

**File Size** The size of the file is shown in bytes. When a fragmented file spans more than one tape or Volume, this column displays the file size stored on the selected tape or Volume followed by the total size of the file in bytes.

**Barcode** This is the name/barcode of the tape(s) where the file is archived.

**Status** The status of the file is displayed as either 'Current', 'Renamed', 'Deleted' or 'Overwritten'.

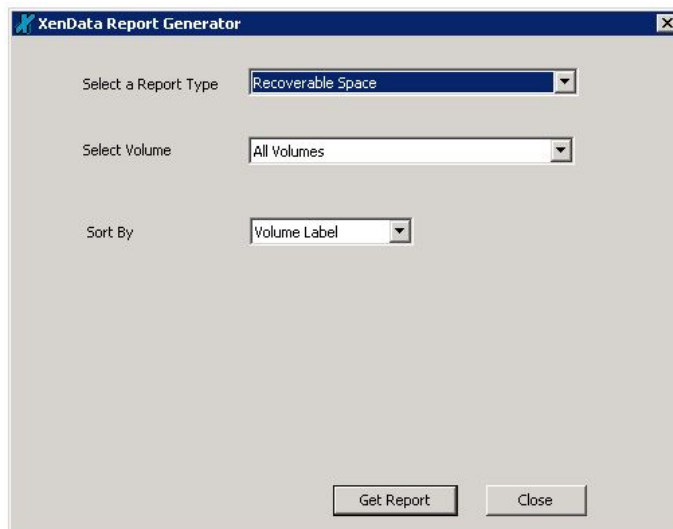
## 10.5 Recoverable Space Reports

### 10.5.1 To Run a Recoverable Space Report

Start the Report Generator; select 'File' and 'New'. Then select the following:

- 'Recoverable Space' as the report type
- 'All Volumes' or a specific Volume
- Sort By either 'Volume Label' or 'Recoverable Space'

This is illustrated below.



Then click 'Get Report'

## 10.5.2 Interpret a Recoverable Space Report

An example of a Recoverable Space Report is shown below. It is especially useful for identifying the amount of space that can be recovered using a Repack operation which recovers tape space used by deleted files and old versions of files. It is described in Section 3.3.15.

No	Volume Label	Barcode	Used Space (bytes)	Available Space (bytes)	Recoverable Space (bytes)
1	4E1C0743-00000001-4EE72290	Barcode:RW0003L4	2042052608	835984080896	64012288
2	4E1C0743-00000003-4FACDB28	Barcode:000022L Barcode:LTO0007BL	2629808128	450833104896	10649600

The display columns are explained below.

**No.** This is the report line number.

**Volume Label** This is the Volume Label for the tape cartridge.

**Barcode** This is the name/barcode of the tapes in the Volume.

**Used Space (bytes)** This is the used space for the Volume.

**Available Space (bytes)** This is the available free space for the Volume.

**Recoverable Space (bytes)** This gives the number of bytes recoverable from the Volume using a repack operation.

## 10.6 Tape Cartridge Contents Reports

### 10.6.1 To Run a Tape Cartridge Contents Report

The dialog box 'XenData Report Generator' contains the following fields and controls:

- Select a Report Type:** A dropdown menu with 'Tape Cartridge Contents' selected.
- Select Tape:** A dropdown menu with 'Barcode:000012L1' selected.
- Select Filter:** A dropdown menu with 'All files' selected.
- Search:** A section containing:
  - Type:** A dropdown menu with 'No Search' selected.
  - Sort by:** A dropdown menu with 'Date' selected.
  - File Name with Text:** An empty text input field.
- Buttons:** 'Get Report' and 'Close' buttons at the bottom right.

The Tape Cartridge Contents Report lists the contents of the selected tape. The displayed report can be filtered to show one of the following:

- **All Files** This displays all files on the tape including deleted files, old versions of files and renamed files.
- **Only Current Files** This displays only the current files. These are files that can be accessed via the archive file system interface and excludes deleted files, old versions of files and renamed files.
- **Only Deleted Files** This displays only deleted files.

You can further filter the results by using a search type other than **No Search**. **File Name Text Search** or **Regular Expression Search** are available search options which will filter the displayed results. When **File Name Text Search** is chosen, the search option supports wildcards.

Having selected the tape and the filtering options, select the **Sort by** option and then click **Get Report**.

**Note:** A Tape Cartridge Contents Report will search only in the Volumes that have a tape contents catalog file cached on the system. Volumes created with versions of XenData Archive Series software prior to 4.50 do not automatically have tape catalogs. Also tapes that have not been finalized and have been imported from another archive system will not automatically have tape catalogs.

### 10.6.2 Interpret a Tape Cartridge Contents Report

An example of a Tape Cartridge Contents Report is shown below.

**XenData Report Generator - [Report5]**  
 File Edit View Window Help

Report Type: Tape Cartridge Contents  
 Tape: Barcode:X0015A  
 Showing: All Files  
 Search Type: None  
 Sorted by: Date

No	File Name	Generation	Version	FileSize (bytes)	Date Archived	Type
1	/Venice/Venice1/943MB_May_2006.avi	0	1	989,532,160	Nov 26 2011 13:51	Current
2	/Venice/Venice1/1_3_8GB-May_2006.avi	0	1	1,487,700,480	Nov 26 2011 13:51	Current
3	/Venice/Venice1/392MB_May_2006.avi	0	1	411,132,928	Nov 26 2011 13:51	Current
4	/Venice/Venice2/1_3_8GB-May_2006.avi	0	1	1,487,700,480	Nov 26 2011 13:53	Current
5	/Venice/Venice2/392MB_May_2006.avi	0	1	411,132,928	Nov 26 2011 13:53	Current
6	/Venice/Venice2/943MB_May_2006.avi	0	1	989,532,160	Nov 26 2011 13:53	Current
7	/Venice/Venice3/943MB_May_2006.avi	0	1	989,532,160	Nov 26 2011 13:55	Deleted
8	/Venice/Venice3/1_3_8GB-May_2006.avi	0	1	1,487,700,480	Nov 26 2011 13:55	Deleted
9	/Venice/Venice3/392MB_May_2006.avi	0	1	411,132,928	Nov 26 2011 13:55	Deleted
10	/Venice/Venice4/1_3_8GB-May_2006.avi	0	1	1,487,700,480	Nov 26 2011 13:55	Overwritten
11	/Venice/Venice4/392MB_May_2006.avi	0	1	411,132,928	Nov 26 2011 13:55	Overwritten
12	/Venice/Venice4/943MB_May_2006.avi	0	1	989,532,160	Nov 26 2011 13:55	Overwritten
13	/Venice/Venice4/943MB_May_2006.avi	0	2	989,532,160	Nov 26 2011 13:56	Current
14	/Venice/Venice4/1_3_8GB-May_2006.avi	0	2	1,487,700,480	Nov 26 2011 13:56	Current
15	/Venice/Venice4/392MB_May_2006.avi	0	2	411,132,928	Nov 26 2011 13:56	Current

The Tape Cartridge Contents Report lists the contents of the selected tape. The display columns are described below.

**No.** This is the number of the file in the display sorted by either date or file name, as defined by the **Sort by** selection.

**File Name** This is the file name including full path from the root of the archive logical drive letter.

**Generation** When a file of a given name and path is first created, it is termed generation 0. Every time the file is deleted or renamed and then a new file of the same name is created, the system increments the generation number. Note that each time the generation number is incremented, the version sequence starts again, with version 1 of the new file being the first that contains data.

**Version** If a file is updated with a newer version by overwriting or appending, the XenData software assigns a new version number. A file's version number increases by one every time it has data written to it. Note that the version number does not increase for every individual write operation, just for every file open that is followed by a write. Version 0 of a file never contains any data; the first time an application writes to the file, the version number is incremented to 1.


**File Size** The size of the file is shown in bytes. When a fragmented file spans more than one tape or Volume, this column displays the file size stored on the selected tape or Volume followed by the total size of the file in bytes.

**Date Archived** Displays the date and time the file is archived.

**Type** The status of the file is displayed as either 'Current', 'Renamed', 'Deleted' or 'Overwritten'.

## 10.7 UnArchived Files Reports

### 10.7.1 To Run an UnArchived Files Report



The UnArchived Files Report lists all the files which are not fully archived to tape that should have been written to tape using the current File Group rules. Select the folder to be searched. Note that all sub-folders will be included in the search.

You can search for a file name with specified text or by using a Regular Expression. When **File Name Text Search** is chosen, the search option supports wildcards.

Having selected the folder and any search option, select the **Sort by** option and then click **Get Report**.

### 10.7.2 Interpret an UnArchived Files Report

An example of an UnArchived Files Report is shown below.

No	File Name	Generation	Version	Replica	Barcode	Status
1	/newtest/001.dat	0	1	1 2	Barcode:000033L Barcode:000102L	Not Archived Archived
2	/newtest/002.dat	0	1	1 2	Barcode:000033L Barcode:000102L	Not Archived Archived
3	/newtest/003.dat	0	1	1 2	Barcode:000033L Barcode:000102L	Not Archived Archived
4	/newtest/004.dat	0	1	1 2	Barcode:000033L Barcode:000102L	Not Archived Archived
5	/newtest/005.dat	0	1	1 2	Barcode:000033L Barcode:000102L	Not Archived Archived
6	/newtest/006.dat	0	1	1 2	Barcode:000033L Barcode:000102L	Not Archived Archived
7	/newtest/007.dat	0	1	1 2	Barcode:000033L Barcode:000102L	Not Archived Archived
8	/newtest/008.dat	0	1	1 2	Barcode:000033L Barcode:000102L	Not Archived Archived
				1	Barcode:000033L	Not Archived

The UnArchived Files Report lists all the files which are not fully archived to tape that should have been written to tape using the current File Group rules. The display columns are described below.

**No.** This is the number of the file in the display sorted by either date or file name, as defined by the **Sort by** selection.

**File Name** This is the file name including full path from the root of the archive logical drive letter.

**Generation** When a file of a given name and path is first created, it is termed generation 0. Every time the file is deleted or renamed and then a new file of the same name is created, the system increments the generation number. Note that each time the generation number is incremented, the version sequence starts again, with version 1 of the new file being the first that contains data.

**Version** If a file is updated with a newer version by overwriting or appending, the XenData software assigns a new version number. A file's version number increases by one every time it has data written to it. Note that the version number does not increase for every individual

write operation, just for every file open that is followed by a write. Version 0 of a file never contains any data; the first time an application writes to the file, the version number is incremented to 1.

**Replica** This column identifies when a file has a File Group rule that indicates it should have been saved to a replicated Volume Set. When a file should not be saved to a replicated Volume Set, the entry in this column is left blank.

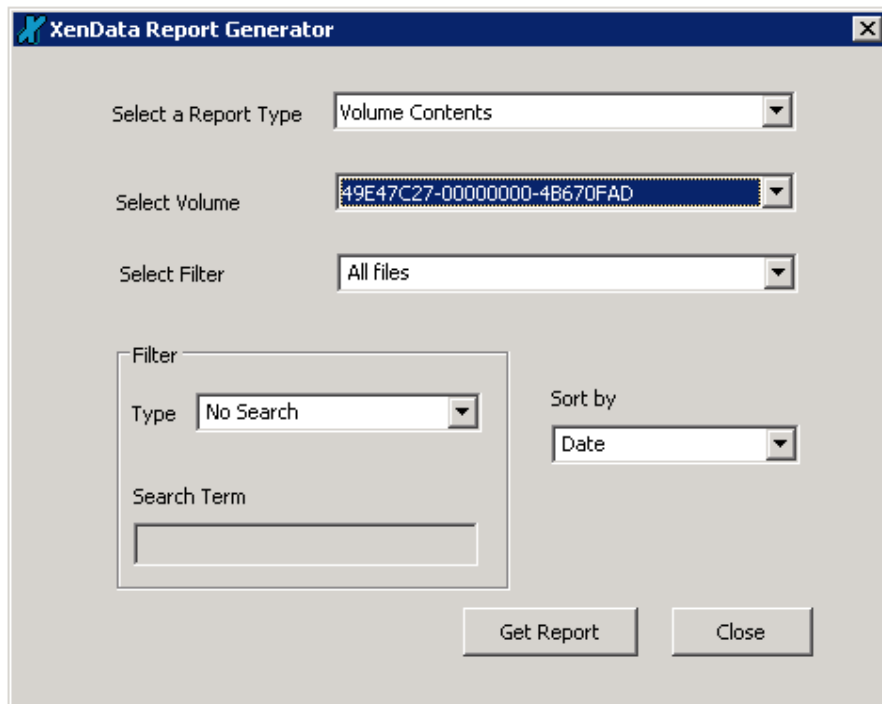
**Barcode** This is the barcode of the tape cartridge on which the file can be archived. This column is left blank when barcode is not available.

**Status** An archived file is listed in this report only when one or more of the replicas are not archived properly. The status of the file instance is displayed as one of the following:

- **Not Archived:** The file is not archived on tape.
- **Partially Archived:** The file is not fully archived.
- **Unverified Archived:** The file may or may not be archived correctly.
- **Archived:** This instance of the file is archived correctly. (You will see this displayed when at least one other instance of the file is not archived correctly to a replica tape.)

## 10.8 Volume Contents Reports

### 10.8.1 To Run a Volume Contents Report



The Volume Contents Report lists the contents of the selected Volume. The displayed report can be filtered to show one of the following:

- **All Files** This displays all files on the Volume including deleted files, old versions of files and renamed files.
- **Only Current Files** This displays only the current files. These are files that can be accessed via the archive file system interface and excludes deleted files, old versions of files and renamed files.
- **Only Deleted Files** This displays only deleted files.



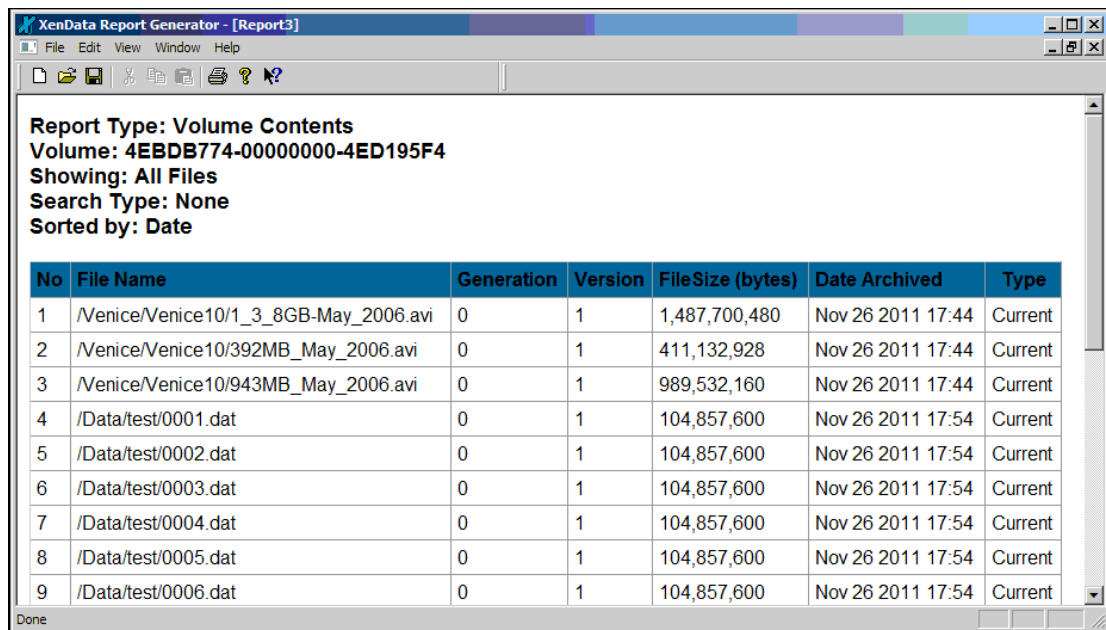
You can further filter the results by using a search type other than **No Search**. **File Name Text Search** or **Regular Expression Search** are available search options which will filter the displayed results. When **File Name Text Search** is chosen, the search option supports wildcards.

Having selected the Volume and the filtering options, select the **Sort by** option and then click **Get Report**.

**Note:** A Volume Contents Report will search only in Volumes that have a tape contents catalog file cached on the system. Volumes created with versions of XenData Archive Series software prior to 4.50 do not automatically have tape catalogs. Also, tapes that have not been finalized and have been imported from another archive system will not automatically have tape catalogs.

## 10.8.2 Interpret a Volume Contents Report

An example of a Volume Contents Report is shown below.



**XenData Report Generator - [Report3]**

File Edit View Window Help

Report Type: Volume Contents  
 Volume: 4EBDB774-00000000-4ED195F4  
 Showing: All Files  
 Search Type: None  
 Sorted by: Date

No	File Name	Generation	Version	FileSize (bytes)	Date Archived	Type
1	/Venice/Venice10/1_3_8GB-May_2006.avi	0	1	1,487,700,480	Nov 26 2011 17:44	Current
2	/Venice/Venice10/392MB_May_2006.avi	0	1	411,132,928	Nov 26 2011 17:44	Current
3	/Venice/Venice10/943MB_May_2006.avi	0	1	989,532,160	Nov 26 2011 17:44	Current
4	/Data/test/0001.dat	0	1	104,857,600	Nov 26 2011 17:54	Current
5	/Data/test/0002.dat	0	1	104,857,600	Nov 26 2011 17:54	Current
6	/Data/test/0003.dat	0	1	104,857,600	Nov 26 2011 17:54	Current
7	/Data/test/0004.dat	0	1	104,857,600	Nov 26 2011 17:54	Current
8	/Data/test/0005.dat	0	1	104,857,600	Nov 26 2011 17:54	Current
9	/Data/test/0006.dat	0	1	104,857,600	Nov 26 2011 17:54	Current

Done

The Volume Contents Report lists the contents of the selected Volume. The display columns are described below.

**No.** This is the number of the file in the display sorted by either date or file name, as defined by the **Sort by** selection.

**File Name** This is the file name including full path from the root of the archive logical drive letter.

**Generation** When a file of a given name and path is first created, it is termed generation 0. Every time the file is deleted or renamed and then a new file of the same name is created, the system increments the generation number. Note that each time the generation number is incremented, the version sequence starts again, with version 1 of the new file being the first that contains data.

**Version** If a file is updated with a newer version by overwriting or appending, the XenData software assigns a new version number. A file's version number increases by one every time it has data written to it. Note that the version number does not increase for every individual write operation, just for every file open that is followed by a write. Version 0 of a file never contains any data; the first time an application writes to the file, the version number is incremented to 1.

**File Size** The size of the file is shown in bytes. When a fragmented file spans more than one tape or Volume, this column displays the file size stored on the selected tape or Volume followed by the total size of the file in bytes.

**Date Archived** Displays the date and time the file is archived.

**Type** The status of the file is displayed as either 'Current', 'Renamed', 'Deleted' or 'Overwritten'.

## Chapter 11. Glossary

Terms used in this manual and other XenData documentation are described in this section.

**Activation Code** An Activation Code is required to run XenData6 Server software and this enables the chosen library and drive configuration.

**Alternate Data Streams** Additional data streams that can be associated with a file name. Also called 'Named Streams' and 'NTFS Streams'.

**API** An acronym for 'Application Program Interface'. Most installations of XenData Archive Series software do not use the XenData API. However, a software development kit is available to developers that allows them to use the XenData API to tightly integrate their application with XenData6 Server software.

**AppleDouble File** A term used by Apple to describe how files may be written to a non-Apple SMB network share. In addition to the main file, a small file containing file attributes is also written. The main file is termed the 'data fork' and the file with attribute data is termed the 'resource fork'. The resource fork file name is pre-pended with the characters '.\_'.

**Archive Series Software** XenData6 Server software is the sixth generation of XenData's Archive Series.

**Autoloader** A term for a tape library which has only one tape drive and cannot be expanded to include more.

**Blank Media Set** The set of tapes shown in the XenData Management Console which consist of new (unused) tapes or rewritable tapes that have been reformatted by the system administrator.

**CIFS** An acronym for 'Common Internet File System', a term promoted by Microsoft. It is the standard protocol used by Windows computers to communicate on a network. It is based on the SMB (Server Message Block) network protocol. Apple OS X documentation often refers to the CIFS/SMB network protocol, simply as 'SMB'.

**Disk Cache** In XenData documentation, the magnetic or solid state disk volume under XenData control is often referred to as the disk cache.

**Dynamic Disks** From Windows 2000 on, Microsoft introduced the option to configure magnetic disk storage as either Dynamic Disks or Basic Disks. The magnetic disk logical drive that is managed by XenData6 Server software should be configured as a Dynamic Disk except when implementing a clustered server arrangement.

**File Fragmentation** The way in which computer systems break large files into smaller, more manageable units for transfer to or from storage devices. Enabling file fragmentation for a File Group allows partial reading of large files and tape cartridge spanning.

**File Group** A File Group is a collection of files that have the same XenData file management policy and consequently are all treated in the same way by the system. Every file that is handled by the system belongs to exactly one File Group. Files are assigned to a File Group on the basis of their names. This assignment can be based on the name of the folder that contains a file, the name of the file or both. When the 'Save files to tape' option is enabled for a File Group, all files in that File Group will be saved to a designated Volume Set.

**Finalizing** Finalizing a tape cartridge written with the TAR format results in special tape marks being written to the tape that prevent additional data files or file fragments being written to that tape.

**Flush** Files are flushed from a storage device when they are deleted to free space. XenData6 Server software can be configured to automatically flush files from disk, once they are securely stored on tape. After flushing, the file is still present in the same location in the file system and is available from tape.

**FTP** An acronym for '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.

**Generation** When a file of a given name and path is first created, it is termed generation 0. Every time the file is deleted or renamed and then a new file of the same name is created, the system increments the generation number. Note that each time the generation number is incremented, the version sequence starts again, with version 1 of the new file being the first that contains data.

**History Explorer** This is enhanced functionality provided within Windows Explorer that is used to obtain the complete history and status of any file which the system knows about. It can list all available versions of all files, including deleted and renamed files. It allows retrieval of all file versions from magnetic disk or tape.

**HSM** An acronym for 'Hierarchical Storage Management'. XenData6 Server software supports three main levels of storage hierarchy:

- **Online** with one instance of a file on the magnetic disk under XenData control and, in addition, there may be one or more instances on tape. In this case the file will be retrieved from disk when it is read.
- **Near-line** with at least one instance of a file on tape within a tape library or tape drive and no instance on magnetic disk. Often the file will be written to more than one tape cartridge using XenData's automatic tape cartridge replication capabilities.
- **Offline** with no instance on magnetic disk and one or more instances of a file on tape, all of which have been exported from the tape libraries and ejected from stand-alone drives. In this case, the file cannot be read until it is put back into a library or a stand-alone tape drive.

**LTFS** An acronym for 'Linear Tape File System'. It is a popular format that defines how file data and file system metadata are written to tape media. It is supported by LTO-5 and later generations of LTO drives.

**LTO** An acronym for 'Linear Tape Open', the most popular mid-range tape format which is also known as Ultrium.

**MMC** An acronym for 'Microsoft Management Console', which is an extensible common presentation service for management applications. MMC can be used to create, save, and open administrative tools that manage the hardware, software and network components of a Windows system. The XenData Management Console is an example of such a tool.

**Named Streams** See Alternate Data Streams

**Near-line** See HSM

**NTFS** Microsoft's advanced file system used to store and manage files on a storage medium. It is the preferred Windows file system when storing files on magnetic disk.

**NTFS Streams** See Alternate Data Streams

**Offline** See HSM

**Offline File Attribute** XenData6 Server software uses the standard Windows offline file attribute to identify files that are no longer online on magnetic disk storage. After a file has been successfully written to tape and is no longer on magnetic disk, the offline attribute is set. This Microsoft file attribute changes network timeout periods to allow retrieval of files from storage media with long access times.

**Online** See HSM

**Petabyte** One thousand (or 1024) terabytes. It is abbreviated to PB.

**Primary Tape** When new files are written to a Volume Set, they are first written to the cartridge that is designated as the Primary. For any given Volume Set, the system may change the physical cartridge that is designated as the Primary. This occurs, for example, when a Primary Cartridge is exported and replaced with one of its replicas. The Primary Cartridge is identified within the XenData Management Console.

**Quarantined Media Set** This is a set of cartridges shown in the XenData Management Console containing tape cartridges that have been imported into the library but for some reason cannot currently be used by the system. Typically, this will be because a cartridge has previously been used by a different application (such as a backup application) or because the tape was repacked using XenData6 Server software.

**RAID** An acronym for 'Redundant Array of Independent Disks'.

**RAM** An acronym for 'Random Access Memory'.

**Repack** Repack is a XenData operation which recovers tape space used by deleted files and old versions of files. The operation may be performed only on Volumes that are not writable, such as are full, finalized and write-protected Volumes.

**SAIT** An acronym for 'Super Advanced Intelligent Tape', a legacy mid-range tape format developed by Sony.

**SMB** See CIFS

**Stand-alone Tape Drive** Data tape drive that is not internal to a tape library.

**Tape Alert** Tape drives and other hardware employ a mechanism called Tape Alert to send maintenance information to external devices. XenData6 software periodically monitors Tape Alert information and informs the user or takes other appropriate action where necessary.

**Tape Catalog** A Tape Catalog contains an index of the contents of a specific Volume. In XenData documentation, it is termed either a 'Tape Contents Catalog' or a 'Tape Catalog' for short. The ability to create a Tape Catalog was introduced in Version 4.50 of XenData Archive Series software.

**TAR** A term derived from 'Tape ARchive'. It is a popular open standard format that defines how file data and file system metadata are written to data tape media.

**Terabyte** One thousand (or 1024) gigabytes. It is abbreviated to TB.

**Ultrium** The most popular mid-range tape format, also known as LTO.

**Volume** If replication is not enabled, the XenData term 'Volume' refers to an individual cartridge. If replication is enabled, a Volume consists of a complete set of replica cartridges which, when up-to-date, all contain the same data. For example, when replication is enabled and one additional copy of each cartridge has been specified by the administrator, a Volume will consist of two tape cartridges.

**Volume Set** This is a set of one or more Volumes which stores files from designated File Groups. As more data is written to a Volume Set, the initial Volume will eventually become full. At a preset threshold, defined by the administrator, the system will automatically add another Volume by taking the appropriate number of tapes from the blank media set and will extend the Volume Set.

**WORM** An acronym for 'Write Once Read Many'. 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-rewritable medium in such a way that files can be 'deleted' within the file system in the normal way. However, they are really only being hidden from the normal file system interface and can be retrieved using the 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.

**XenData Alert Module** This is a companion product which is installed separately. It provides e-mail and on-screen alerts that are tailored to the needs of archive system operators, systems administrators and IT support personnel. The alerts are derived by filtering and categorizing events recorded in the XenData Event Log.

**XenData Event Log** This is part of Windows Event Viewer within the Computer Management Console. The XenData Event Log is used for XenData specific messages. If the system is not behaving as expected, the Event Log is the first place that a system administrator should look.

**XenData Management Console** This is used to configure all File Group, Volume Set and tape replication options, and to view diagnostic information about the system.

**XenData State File** This contains tape information and the XenData Management Console settings, including File Group and Volume Set configuration settings.