

XenData6 Server Software

USER MANUAL

Version 6.22

© 2003-2017 XenData Limited. All rights reserved.

1. Overview	6
1.1 Hierarchical Storage Management	8
1.2 Windows Offline Attribute	9
1.3 Tape Cartridge Formats	10
1.4 Optical Disc Cartridge Format	10
1.5 Robotic Library and Stand-Alone Drive Configurations	11
1.6 Renaming Folders	12
1.7 Barcode Management	12
1.8 Cache disk requirements	13
1.9 Compatible Operating Systems	14
1.10 Antivirus Software Compatibility	14
1.11 License Administration	14
2. Concepts	16
2.1 Writing Files to the Archive	17
2.2 Reading Files from the Archive	18
2.3 File Groups	18
2.4 Volumes, Volume Sets and Automatic Replication	19
2.5 Partial File Restore and Cartridge Spanning	20
2.6 Pending Write Mode	20
2.7 Offline File Management	21
2.8 File Version Management	21
2.9 File Security	22
2.10 Handling of File Delete and Rename Operations	22
2.11 Volume Contents Catalogs	22
2.12 Volume Finalization	23
2.13 Recovering Cartridge Space using Repack	23
2.14 Alternate Data Streams	24
3. Administering the System	26
3.1 New Installation Check List	28
3.1.1 Special Considerations for Creating Directories at the Root of the Archive	29
3.2 File Groups	29
3.2.1 Allocating Files to a File Group	30
3.2.1.1 Creating a New File Group	30
3.2.1.2 Renaming a File Group	31
3.2.1.3 Changing the Order of File Groups	32
3.2.1.4 Examples of Allocating Files to File Groups	32
3.2.2 Setting Policies for File Groups	34
3.2.2.1 How to Construct the "File Name or Path Pattern" or the "Exclude pattern" for a File Group	34
3.2.2.2 Selecting Storage Options for a File Group	35
3.2.2.3 File Fragmentation	35
3.2.2.4 Disk Retention Rules	36

3.2.2.5	Changing Retention Rules	38
3.2.2.6	File Group Advanced Options	38
3.3	Volume Sets	41
3.3.1	The Blank Cartridge Set	42
3.3.2	The Quarantined Cartridge Set	42
3.3.3	Volume Set Defaults for a New Installation	43
3.3.4	Configuring a Volume Set	44
3.3.5	Configuring Replication for a Volume Set	45
3.3.6	Defining Replication Timing	46
3.3.7	Allocating Volumes to a Volume Set	47
3.3.8	Creating a New Volume Set	48
3.3.9	Removing Information about a Cartridge from the System	49
3.3.10	Replacing a Missing Replica Cartridge	50
3.3.11	Renaming a Volume Set	51
3.3.12	Deleting a Volume Set	52
3.4	Managing Data Cartridges	53
3.4.1	About Rewritable Cartridges	53
3.4.2	About WORM Cartridges	53
3.4.3	Importing Cartridges	54
3.4.4	Transferring Cartridges Between Systems	54
3.4.4.1	Considerations for LTFS	55
3.4.4.2	Considerations for ODA	55
3.4.4.3	Considerations for TAR	56
3.4.5	Exporting Cartridges	56
3.4.6	Reformatting Cartridges	57
3.4.7	Displaying Cartridge Information	59
3.4.8	Volume Statistics	61
3.4.9	Write-Protecting Cartridges	63
3.4.10	Building and Rebuilding Volume Contents Catalogs	65
3.4.11	Repacking Volumes	67
3.4.12	Canceling a Repack Operation	69
3.4.13	Verifying the Data on a Cartridge	70
3.4.14	Finalizing Volumes	71
3.4.15	Setting the Administrator Defined Information for a Cartridge	72
4.	Using Mac Clients	74
4.1	Hidden File Group Policies	75
4.2	Disabling Icon Preview	76
4.3	Disabling Alternate Data Streams	76
4.4	Support of OS X Characters	77
5.	Windows Explorer Extensions	78
5.1	Smart Copy and Paste	79

5.2	Flushing	80
5.3	Pre-fetching	81
5.4	Volume View	83
5.5	History Explorer	83
5.6	Enhanced Properties for the Archive Volume	85
6.	Metadata Backup	88
6.1	Starting Metadata Backup	89
6.2	Selecting Backup or Restore	89
6.3	Making a Predefined Backup	90
6.4	Making a Custom Backup	93
6.5	Restoring from a Backup	97
7.	Scheduler	101
7.1	Starting the Scheduler	102
7.2	Adding a Task	102
7.3	The Scheduler Status Display	103
7.4	Editing and Deleting Tasks	104
7.5	Starting and Stopping Tasks	104
7.6	Scheduling Metadata Backup	105
7.7	Scheduling Deferred Write	106
8.	System Recovery	107
8.1	Rebuilding a System from Data Cartridges	108
8.2	In Case of Hardware Failure	109
8.2.1	Options in Case of Library or Drive Failure	110
8.2.2	Temporarily Disabling Hardware	110
9.	Diagnostics & Maintenance	112
9.1	Windows Event Log	113
9.2	System Trace Log	114
9.3	Library and Drive Diagnostic Information	115
9.4	Cleaning Tape Drives	116
9.5	Displaying Information about Cleaning Cartridges	117
9.6	Cartridge Alert State	118
10.	Reports	120
10.1	Starting the Report Generator	121
10.2	Creating, Saving and Restoring Reports	121
10.3	File Search Report	123
10.3.1	Interpreting a File Search Report	124
10.4	Recoverable Space Report	125
10.4.1	Interpreting a Recoverable Space Report	126

10.5	Data Cartridge Contents Report	126
10.5.1	Interpreting a Cartridge Contents Report	127
10.6	UnArchived Files Report	128
10.6.1	Interpreting an UnArchived Files Report	129
10.7	Volume Contents Report	131
10.7.1	Interpreting a Volume Contents Report	132
11.	Alert Module	133
11.1	Installation	134
11.2	License Administration	134
11.3	Language Support	135
11.4	Event Monitor About	135
11.5	Event Monitor Configuration	136
11.6	About Event Categories	137
11.7	Configuring Event Categories	137
11.8	About Groups	138
11.9	Configuring Groups	139
11.10	About Email Server	140
11.11	Configuring Email Server	140
11.12	Error Reporting	142
11.13	On-Screen Messaging About	143
11.14	On-Screen Messaging Configuration	143
12.	Glossary	146
	Index	150

Overview

1 Overview

XenData6 Server™ software manages a Windows file server and provides high performance archiving to magnetic or solid state disks, optical discs and data tape cartridges. It uses a magnetic or solid state disk volume for caching and for retention of frequently used content. The software manages standalone drives, robotic libraries and the cache disk, creating a highly scalable digital archive.

XenData6 Server provides the following key functionality:

- **Standard File System Interface** The archive appears as a Windows file system with a single logical drive letter. This represents the files stored on the cache disk, on data cartridges in the attached robotic libraries and stand-alone drives and on data cartridges that have been exported (externalized). Either the entire archive or one or more folders within it may be configured as a network share. The use of a Windows file system means that the archive integrates seamlessly with standard applications and with existing network infrastructures.
- **Hierarchical Storage Management** Administrator-defined policies determine where files are physically stored. Three tiers of HSM are supported: online, nearline and offline. The standard Windows offline file attribute bit identifies when a file has been flushed from the cache disk and is only available from one or more data cartridges.
- **Disk Retention Determined by Policy** The cache retention time is determined by administrator-defined policies.
- **Standard Network Protocols** The solution supports CIFS/SMB and FTP file transfers.
- **Choice of Cartridge Formats** The archived files are written to optical discs using the UDF format or to data tapes using either the open standard TAR format or the LTFS interchange format.
- **Support for Multiple Cartridge Groups** The software allows groups of files to be allocated to specified groups of data 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 tape cartridges that may be exported from the library for off-site retention. Furthermore, the cartridges may be rapidly imported into a replica Disaster Recovery system.
- **Intelligent Bar code Management** The software selects data cartridges in barcode order and also automatically selects matched barcodes for replicated volumes.
- **Repack of Cartridges** This operation copies currently available files (excluding deleted files and old versions of files) to new data cartridges. After completion, the original data cartridges are quarantined and may be reused. Repack is useful for recovery of cartridge capacity lost to deleted files and old file versions. It is also useful for migrating files to different types of cartridges, for example when a new generation of drive supports higher capacity 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.
- **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. Restore requests are processed in an order that minimizes unnecessary cartridge swaps and ensures that multiple files that are read from an individual cartridge are read in the optimum transfer order.
- **Self-Describing Data Cartridge Formats** Every data cartridge contains all the file system metadata necessary to recover all the files stored on it, whichever cartridge format is used.
- **Transfer of Content between Systems** Export and import functions make it easy to transfer cartridges from one system to another.
- **Microsoft Security** The archive file system is fully integrated with the Microsoft Windows security model based on Active Directory.

- **WORM Support** The software supports tape and optical disc cartridges using standard rewritable and unalterable WORM cartridge formats.
- **File Version Control** The software provides comprehensive file version control. Deleted files and old file versions may be restored from the archive (unless the files have been purged using the repack operation).

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

- **Metadata Backup** - The archive metadata backup program backs up and restores the archive file system metadata and the Archive Management Console settings. It allows rapid restore of the system in case of failure of the cache disk.
- **Report Generator** - The report generator creates archive system reports including a report showing the contents of any data cartridge.
- **Scheduler** - The Scheduler can be used to schedule metadata backups and to defer writing to data cartridges to a scheduled period.

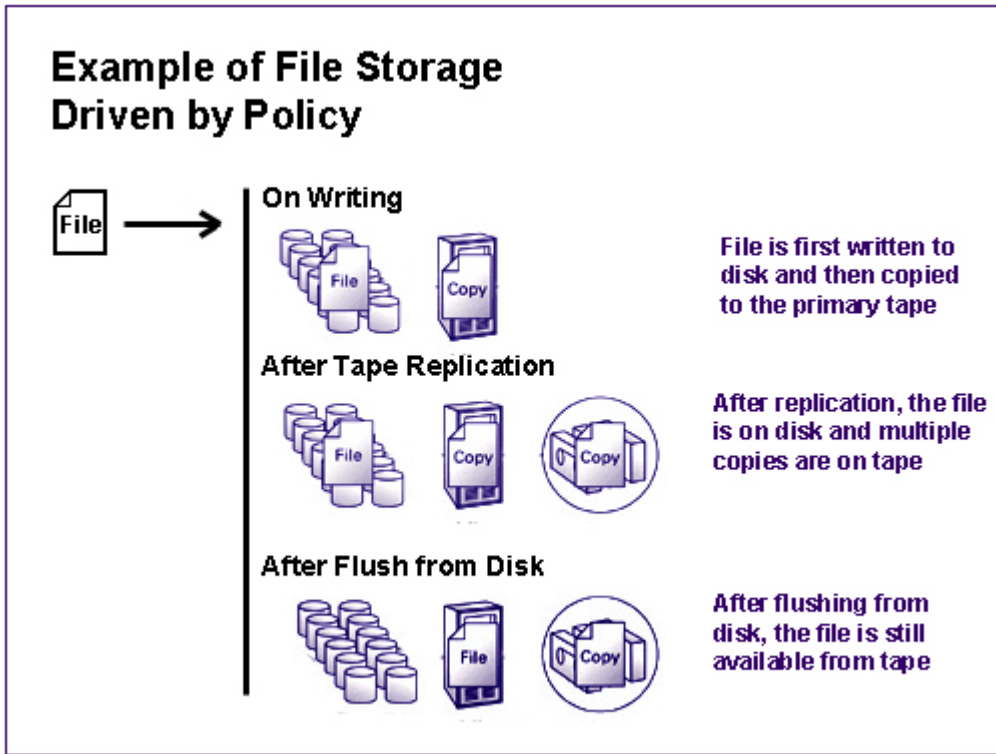
The Alert Module, which is installed separately, provides email alerts and on-screen notifications that are tailored to the needs of archive system operators, system administrators and IT support personnel. The alerts are derived by filtering and categorizing events recorded in the Windows Event Log.

1.1 Hierarchical Storage Management

XenData6 Server software supports three levels of hierarchical storage:

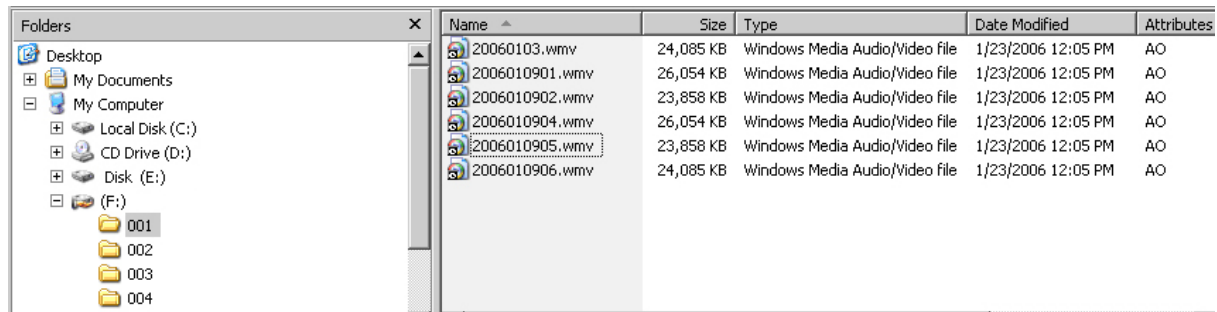
- **Online** with one instance of a file on magnetic or solid state disk (there may also be additional instances on data cartridges). Files that are read when they are in this state will be restored from the cache disk.
- **Nearline** with at least one instance of a file on a data cartridge located within a tape or optical disc drive or library and no instance on the cache disk. Often the file will be written to more than one data cartridge using XenData6 Server's automatic cartridge replication capabilities. Files that are read when they are in this state will be restored from a data cartridge and retained for a predefined time on the cache disk.
- **Offline** with no instance on the cache disk and one or more instances on data cartridges, all of which have been exported from the available drives and libraries. An attempt to read a file that is in this state will fail and a message will be delivered informing the system administrator of the identities of the cartridges that contain the data.

The administrator establishes [cache retention rules](#) that determine when different kinds of files are to be removed from the disk cache. Once the rules have been established, the software manages everything automatically. A typical storage policy is illustrated below.



1.2 Windows Offline Attribute

XenData6 Server software uses the standard Windows offline file attribute to identify files that are no longer present online on the cache disk. After a file has been successfully written to data cartridges 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 some versions of Windows (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:

HSM tier	Description	Offline attribute
Online	One instance of the file is on disk and, in addition, there may be one or more instances on data cartridges.	Not set
Near-line	At least one instance of the file is on a data cartridge within a	Set

Offline robotic library or stand-alone drive and no instance is on disk. There is no instance of the file on disk and there are one or more Set instances on data cartridges, all of which have been exported from robotic libraries or ejected from stand-alone drives.

1.3 Tape Cartridge Formats

All versions of the software may be licensed to support the TAR (Tape Archive) tape cartridge file system format and XenData6 Server versions 6.10 and above may be licensed to support the LTFS (Linear Tape File System) format. These formats define how data is written to the tape cartridge. 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 to use data tapes, the administrator selects the tape cartridge type and then either LTFS or TAR as the cartridge file system format.

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 rewritable and WORM tapes. A tape cartridge written using the TAR format is fully self-describing. However, when using a basic TAR implementation, the whole tape must be scanned to determine the tape's contents. XenData6 Server software extends the TAR format by maintaining a Volume Contents Catalog that includes an index of the cartridge contents. When the tape is full, the Volume Contents Catalog is automatically written to the end of the tape in a "finalization" operation. Finalization can also be performed [manually](#). When a finalized tape cartridge written using the TAR format is moved to another XenData6 Server system, the contents are quickly determined because the system automatically reads the catalog from the tape.

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 tape index and a large data partition for the file data. It is applicable to rewritable LTO-5 and later generations of LTO cartridges and the StorageTek T10000C and later generations. A tape cartridge written using the LTFS format is self-describing and the contents of the cartridge can be determined quickly by reading the index partition on the tape.

XenData6 Server supports the following tape drive and cartridge combinations with LTFS:

- LTO-7 drives and 6 TB LTO-7 rewritable cartridges
- LTO-7 drives and 2.5 TB LTO-6 rewritable cartridges
- LTO-7 drives and 1.5 TB LTO-5 rewritable cartridges in read-only mode
- 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
- STK T10000D drives and 8.5 TB rewritable cartridges

1.4 Optical Disc Cartridge Format

XenData6 Server versions 6.20 and above may be licensed to support the UDF optical disc file system format recorded on Sony ODA cartridges. When configuring a Volume Set to use optical discs, the administrator selects the Optical disc cartridge type.

Sony ODA offers an alternative to data tape for the longer-term archiving of valuable assets. The format is designed to be fully backward-read compatible, providing read capability for all generations of ODA cartridges and eliminating the need for cartridge migration due drive obsolescence. This

version of XenData6 Server supports Sony ODA cartridges in 300 GB, 600 GB, 1.2 TB and 1.5 TB capacity sizes, in re-writable and write-once formats.

1.5 Robotic Library and Stand-Alone Drive Configurations

XenData6 Server software may be configured and licensed to manage one or more robotic libraries, optionally combined with one or more stand-alone drives. A Windows file server with a tape library is illustrated in the diagram opposite. A robotic 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 files in the archive appear within one Windows logical drive letter on the server whether those files are on the disk cache, on a data cartridge Volume within a library or are on an offline cartridge. 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 drives within the library is an important consideration and a major factor in the library cost. Factors to consider are as follows.

- For archives with only one 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 Volume replication is scheduled to occur at a time when there is no other writing or reading activity.
- A system with two 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 drives are recommended for large or frequently accessed archives. XenData6 Server software will intelligently manage many drives, allowing simultaneous writing and multiple file read accesses.
- All installations that include a library, even if it has only one drive within it, support cartridge replication.

In a configuration with a library and at least one stand-alone drive, the drive can be very useful for bringing offline cartridges quickly near-line. For installations that do not have a library, only one or more stand-alone drives, setting the Pending Write Mode option is often desirable because it allows the cache disk to act as a buffer when writing a set of files to multiple cartridges.

Regardless of the number of stand-alone drives or robotic libraries in a XenData6 Server system, all files in the archive appear within one Windows logical drive letter on the server whether those files are on the disk cache, on a data cartridge Volume within a stand-alone drive or are on an offline cartridge.



1.6 Renaming Folders

XenData6 Server software does not support renaming folders that contain files.

If a user renames a folder, the fully qualified name (complete file path name) of every file in the folder and every sub-folder changes. In this situation, the audit model used by XenData6 Server software would have to record a separate record for each file in the folder hierarchy below the renamed folder, recording the fact the path name had changed. Because there is no limit on the number of files a sub-folder hierarchy can contain, there is no limit to the length of time this operation could take. In order to avoid uncertainty and provide a consistent interface, XenData6 Server software prohibits changing a folder's name if the folder contains any files. Any attempt to rename a folder in these circumstances will result in a "folder is not empty" error.

Folders that do not contain any data, such as folders that have just been created via the New Folder command in Windows Explorer, can be renamed.

1.7 Barcode Management

Barcode labels are available for all data cartridge formats supported by XenData6 Server software and most robotic libraries include a barcode reader as standard. Barcodes are readable by both humans and robotic libraries as shown below. Barcodes are strongly recommended for keeping track of cartridges when using a robotic library.

Most data cartridge formats, including LTO and ODA, include an in-cartridge memory chip. When a cartridge is used in a robotic library, XenData6 Server software writes the barcode information to the cartridge memory chip. This is particularly useful when the cartridge has been exported from the library and inserted into a stand-alone drive. Stand-alone drives do not contain barcode readers but can read the contents of the cartridge memory. This allows XenData6 Server software to provide a consistent identification of barcode for all data cartridges that have been in a library, even when a cartridge is being used in a stand-alone drive.



Barcodes are used to identify data cartridges in the [Archive Management Console](#), [Event Log](#), [History Explorer](#) and [Report Generators](#). In addition to these functions, XenData6 Server software matches barcodes for replicated tapes and selects tapes in barcode order. The inventory of blank cartridges provided by the library is sorted into alphanumeric barcode order. When allocating tapes for replicated volumes, the system will look for a matched set of barcodes which differ by only one letter (e.g. 'A' and 'B') in one character position. Where possible, the lowest matched set is allocated for replicated sets of tapes. For non-replicated volumes or when no matched set exists, cartridges are allocated in alphanumeric order.

For example, if we had the following sequence of barcodes in the Blank Cartridge Set.

- X0007AL6
- X0008AL6
- X0008BL6
- X0009AL6
- X0009BL6

The system would next allocate X0008AL6 and X0008BL6 to a replicated set of tapes. However, if the system were allocating a cartridge to a non-replicated set, it would select X0007AL6.

Most robotic libraries use a barcode format called '3 of 9' or code 39. The last two digits of the human readable format often represent the data cartridge format. For example L6 is used for LTO-6 tape cartridges. A check digit may or may not be present in the machine readable barcode; this is used to verify the integrity of the other digits in the label. Many libraries can be configured to read barcodes either with or without a check digit. In these cases, it is best to configure the library not to include the check digits as this will give the most consistency in the use of barcodes.

Note that cleaning cartridges have a specific barcode label format which always starts with CLN. This allows XenData6 Server software to recognize a cleaning cartridge without putting it in a drive and unnecessarily using a cleaning cycle.

1.8 Cache 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 XenData6 Server software. The archive management software is installed on the boot drive but it requires minimal available capacity. The dedicated cache disk 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 should be configured as a dynamic disk except in clustered environments.

Capacity requirements for the cache disk 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 cache disk 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 cache disk provides caching for each file that is being written and all files that are open while being read by programs. 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 Archive Management Console to define policies for online retention. If 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 the sizes of all the files.

Example Calculation for an installation with only stand-alone 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 cache disk capacity of 102 GB or higher is required for this installation.

Example Calculation for a 200 TB digital video archive using a robotic 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
- High resolution file size of 100 GB - 2 TB 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 12 TB is estimated for this installation

1.9 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

1.10 Antivirus Software Compatibility

When installing anti-virus protection on an archive running XenData6 Server, it is important to choose an anti-virus (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.

For more information about certified AV products, please refer to the XenData Technical Note XTN1201.

1.11 License Administration

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

To start the License Administration Program

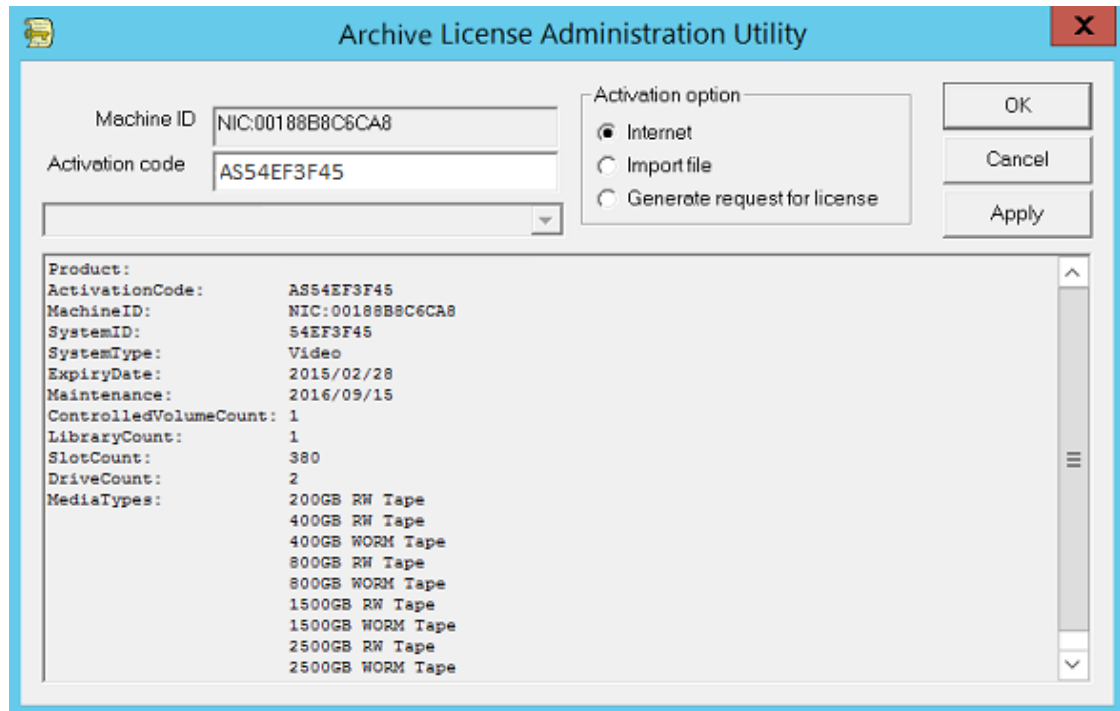
For Windows Server 2008 R2 and Windows 7:

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. Type "Archive License Administration"
3. Click on the **License Administration** entry in the list to the right of the screen.

The license administration program detects the unique network interface card (NIC) ID of the system and prompts for an Activation Code. A digitally signed activation file must then be obtained from a cloud-based license server. There are two automated ways to obtain this via the Internet as described in the first two options below. If the XenData6 Server has no Internet connectivity, the digitally signed activation file can be obtained from 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 a 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, contact your technical support for advice. They will need the Activation Code and the NIC identity information contained in the file created using the **Generate request for license** as described in Option 2 (this file is in HTML format and is human-readable). They will return an XML file that contains the digitally signed license. 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.

Concepts

2 Concepts

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 data cartridges (termed 'Volume Sets');
- groups of files to be retained online on disk for predefined periods of time;
- temporary storage of files on disk in case of no writable data cartridges being available;
- define archive storage format (TAR or LTFS tapes; ODA optical disc cartridges)
- automatic generation of 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.

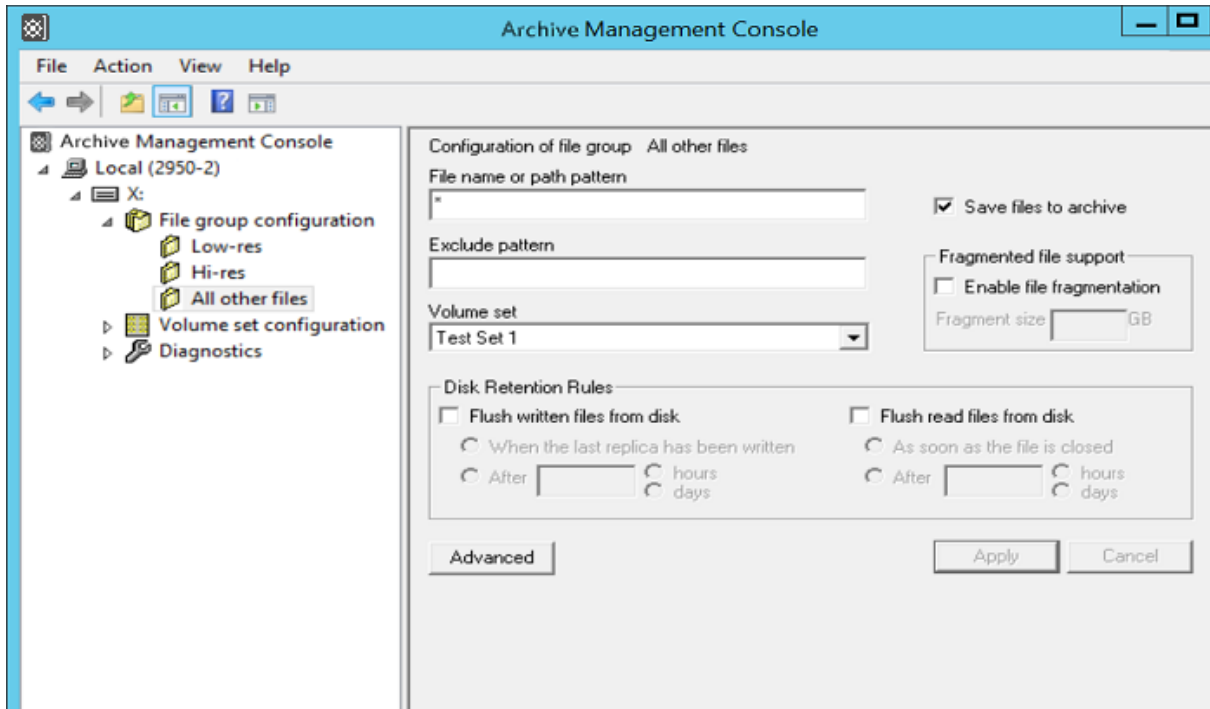
In addition, the following topics are discussed:

- the method used to keep track of the contents of individual cartridges using Volume Contents Catalogs; and
- the repack operation which recovers space occupied by deleted files and old versions of files that are no longer accessible

2.1 Writing Files to the Archive

To write a file to a data cartridge, the file must be [allocated](#) to a File Group and the File Group must have the "[Save files to archive](#)" option selected. The File Group will have an assigned Volume Set which must have one or more Volumes with sufficient free space that are accessible within a robotic library or stand-alone drive.

When these prerequisites have been met, files can be written directly to the system via the standard Windows file system interface.



2.2 Reading Files from the Archive

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

2.3 File Groups

A File Group is a group 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 a ".tmp" extension of folder and these are stored to the cache disk but are not saved to the archive.
- File Group 2 contains all non-temporary files saved to the folder \project1\ which are permanently retained on the cache disk and also saved to a specific Volume Set 'Volumes01'
- File Group 3 contains all non-temporary files saved to the folder \project2\ which are saved to the "Volumes02" Volume Set, retained on the cache disk for 24 hours after they were written or for 48 hours after being read.

Every file and every folder must belong to a File Group; the system will block creation or reading of files for which there is no File Group. In the example above, files can only be written to the folders \project1\ or \project2\ or if they have a ".tmp" extension. Attempts to write other files will not be permitted. After installation of the software, the system is configured with a Default ('catch all') File

Group that sets a policy for all files.

The following parameters are 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 Archive Management Console display, defines the files that are allocated to a particular File Group.
- The position of the File Group relative to other File Groups in the Archive Management Console display, which defines the order in which file name or path pattern matching is applied. If a file name matches the selection criteria for more than one File Group, the file will be assigned to the first matching File Group.
- Enable or disable 'Save files to archive'.
- Selection of a Volume Set, if the File Group is saved to the archive.
- 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 the cache disk, if the File Group is saved to the archive.

The policies defined by the administrator determine how files are stored on the cache disk and data cartridges. The policies do not change the appearance of 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.

2.4 Volumes, Volume Sets and Automatic Replication

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

A Volume Set comprises a set of Volumes that store 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 cartridges from the [Blank Cartridge Set](#) to create a new Volume and extend the Volume Set.

If replication is not enabled then each file is archived to only a single data cartridge. Alternatively, if replication *is* enabled then replica copies of each data cartridge are automatically generated and kept up to date according to a replication schedule. The data on replicated cartridges in a Volume Set are kept synchronized in accordance with the [replication schedule](#) whenever the replica cartridges are available to the system. If one or more replicas are removed from the library, the [Volume Contents Catalog](#) maintains a record of which files need to be written to those cartridges to bring them up to date. When cartridges are reintroduced into the library, the system automatically updates in accordance with the replication schedule.

With the exception of the Blank Cartridge Set and the [Quarantined Cartridge Set](#), all data cartridges within a Volume Set must be of the same type, either Tape or Optical and either WORM or rewritable. Additionally, replicas must all be of the same capacity.

One special Volume Set, termed the [Blank Cartridge Set](#), contains all the cartridges that are present by the system but are not formatted for storing data. These may be new (unused) data cartridges or rewritable cartridges that have been reformatted using the Archive Management Console. Another special Volume Set, termed the [Quarantined Cartridge 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 previously been used by a different application (such as a backup application) and it is recorded in an format that is not compatible with the XenData6 Server software; or
- an error occurred while the system was trying to identify the contents of the cartridge.

After initial configuration of the File Groups, Volume Sets and any associated replication requirements, the system operates completely automatically. Files written to the logical drive under XenData6 Server control are automatically allocated to File Groups. Files allocated to File Groups with 'Save files to archive' enabled have an assigned Volume Set and are automatically written to both disk and data cartridges. If replication is enabled for the assigned Volume Set, this occurs automatically.

Note: XenData6 Server systems support Volume replication with tape based systems. ODA (optical disc) based archives do not support Volume replication because the defect management system used by ODA is not compatible with the XenData6 Server Volume-based replication model. Replication can be enabled for a tape based system that contains a robotic library, even if it has only one drive, and for systems with two or more stand-alone drives. Replication is not supported in a system that has only one stand-alone drive.

2.5 Partial File Restore and Cartridge Spanning

It is often useful in professional video applications to restore a portion of a file without fetching the whole file back to the cache disk. For example, when a short clip is being read from a very large video file, it might take many minutes to restore the whole file. The ability to restore the parts of the file that are needed is called Partial File Restore and it can greatly improve the performance of the system.

Certain application areas, such as Oil and Gas Exploration, generate extremely large files that are bigger than the capacity of current data cartridges. In these applications it is useful to be able to span individual files across multiple cartridges.

Partial File Restore and Cartridge Spanning capabilities are enabled in a XenData6 Server system by using [File Fragmentation](#). The term File Fragmentation refers to the way in which computer systems break large files into smaller, more manageable units for transfer to or from storage devices. File Fragmentation is an optional feature that can be enabled on tape-based XenData6 Server systems. It is usually only worthwhile for files that have a size of several tens of gigabytes or more. File fragmentation is not supported on ODA cartridges because it would break compatibility with the Sony tools and utilities.

2.6 Pending Write Mode

In normal operation, XenData6 Server software writes files to the archive immediately after they have been written to the cache disk. Volume Set configuration settings determine the system's response to an attempt to write files. If this is not possible it is because all of the Volumes in a Volume Set have become full or unavailable. This might happen for a number of reasons:

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

If the Volume Set **Write to disk if no writable volumes are available** option has been enabled and

all Volumes in a Volume Set become full or unavailable, the system automatically enters the Pending Write Mode and will accept more data. However, if this 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 the archive and continues writing to the cache disk. When a writable Volume becomes available within the Volume Set, the system automatically 'catches up' and writes the deferred files to the new Volume. 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 cache disk. When the 'Write to disk if no writable volumes are available' option is enabled, we recommend that the Alert Module be configured to provide notification via email and/or on-screen message of these warning messages.

2.7 Offline File Management

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

- **Online** with one instance of a file on magnetic or solid state cache disk (there may also be additional instances on data cartridges). Files that are read when they are in this state will be restored from the cache disk.
- **Nearline** with at least one instance of a file on a data cartridge located within a tape or optical disc drive or library and no instance on the cache disk. Often the file will be written to more than one data cartridge using XenData6 Server's automatic cartridge replication capabilities. Files that are read when they are in this state will be restored from a data cartridge and retained for a predefined time on the cache disk.
- **Offline** with no instance on the cache disk and one or more instances on data cartridges, all of which have been exported from the available drives and libraries. An attempt to read a file that is in this state will fail and a message will be delivered informing the system administrator of the identities of the cartridges that contain the data.

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 Windows Event Log that identifies which data cartridges contain the file. When the Alert Module is installed, on-screen messages and e-mail alerts are also generated that identify the file name and the cartridges that contain the file. The Alert Module is a companion product that is licensed separately.

2.8 File Version Management

The default Windows file system interface provides access to the latest version of a file but does not permit access to old file versions or to deleted files. XenData6 Server software maintains a complete version history of files in the archive. Old file versions or deleted files can be viewed and restored using [History Explorer](#), which is extended functionality within Windows Explorer that is provided by XenData6 Server software.

When a file is initially created it has a version number of 0. Version 0 does not contain any data; it always has zero size. When an application writes the first byte of data to a file, the version number is incremented to 1. When the file is closed following a version number increase, the file is archived to one or more data cartridges (if the "Save files to archive" option is selected). If the file is subsequently re-opened and has more data written to it, the version number will once again be incremented.

If a file is renamed or deleted and then a new file of the same name is created, the system starts

again with version 0 of the new file and a new generation is created. The generation number starts at 0 the first time the file is created and increases by one every time a new file is created.

Example: the very first time data is written to a new file, it will have generation 0 and version 1. If the file is then reopened and has more data appended then it will be at generation 0 version 2. If the same file is deleted and then data is written to a new file with the same name, the new file will have generation 1 and version 1.

2.9 File Security

XenData6 Server software integrates fully with the Microsoft Windows security model, based on Active Directory. Files and folders 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 any other 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.

2.10 Handling of File Delete and Rename Operations

A change was made in version 6.11 of XenData6 Server software in the way delete and rename operations are handled. This change was made to ensure that every data cartridge is completely self-contained even when files are deleted or renamed after the cartridge becomes full,

In XenData6 Server versions prior to 6.11, a delete or rename record was written to the current writable Volume (i.e. the Volume where new files would be written) which is not necessarily the volume containing the file. From version 6.11 these records are written to the volume that contains the file. Therefore the applicable data cartridge must be available for writing, i.e. not write protected and mounted in a stand-alone drive or in a robotic library. It is not possible to rename or delete a file if the cartridge that contains the file is not available to write. A message in the Event Log states that the required cartridge is offline and gives the cartridge barcode label. When the cartridge is made available nearline, 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 available for the delete or rename to be successful. In this case, any offline replicas will be identified in the Archive Management Console as 'Needs updating', and will be updated when 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 applicable record will be written to all cartridges that contain that file and all the cartridges must be available.

Note: It is possible to override this behavior for file deletes by setting the [File Group Advanced Option](#) "Do not preserve history for deleted files".

2.11 Volume Contents Catalogs

A Volume Contents Catalog contains a journal of the history of a Volume and includes an index of the files and folders on the volume. In this documentation, it is termed either a 'Volume Contents Catalog' or a 'Volume Catalog' for short. The presence of a Volume Catalog is not always required to write,

read or access files but it is required for successful completion of certain other management functions including Repack and to generate a report of cartridge or Volume contents. Furthermore, it greatly reduces the time to perform an Import Folder Structure operation, which is an advantage when transferring cartridges from one XenData6 Server archive to another and when rebuilding the file system on the cache disk volume.

When a new Volume is initially created and added to a Volume Set, the system creates a Catalog in a hidden folder on the cache disk. As folders and files are added and perhaps renamed or deleted, the Volume Catalog is updated. When a XenData6 Server system imports an unknown (or updated) cartridge that contains data written on another system, it will attempt to build a Volume Catalog. In the case of [Finalized](#) TAR format tapes the Catalog is read from the end of the tape. In the case of non-Finalized TAR tapes the system does not attempt to build a Catalog because this operation involves reading the entire tape and might take several hours; the administrator has the option of performing this function if required. In the case of ODA cartridges and LTFS format tapes the Volume Catalog is built from the most recent version of the index data structure recorded on the cartridge. This is done both for cartridges that were written on a XenData6 Server system and for cartridges that were written by other implementations.

When a Volume Catalog has been built from the index on an LTFS or ODA cartridge, the Catalog does not contain information about older versions of the index that may include files which have subsequently been renamed or deleted. If a complete Volume Contents Catalog is required (for example, to recover deleted files or to ensure that the system can account for every byte of data on the cartridge) then the [Rebuild Catalog](#) operation should be used. In cases where a Volume Catalog does not exist on the cache disk (for example, import of non-Finalized TAR format tape) the [Build Catalog](#) operation can be used.

2.12 Volume Finalization

Volume finalization is only applicable to TAR formatted tapes and ODA WORM cartridges. It is a process that writes a special sequence to a data cartridge to indicate the end of the recorded data. In the case of TAR formatted tapes this special sequence is two 512-byte blocks of zeros (a TAR "end of archive" record). In the case of ODA WORM cartridges finalization closes the recording session on the optical discs in the cartridge.

In the case of TAR formatted data tape cartridges the end of archive mark is followed by the Volume Catalog and a file mark. Because they follow the end of archive marker, these items are invisible to standard TAR readers. This is a XenData6 Server proprietary extension to the TAR format that optimizes access to the contents of the tape by putting a complete tape index in a known, easily accessible location (the end of the tape). This optimization is particularly advantageous when transferring tapes from one XenData6 Server system to another or when rebuilding a system from the tape cartridges.

2.13 Recovering Cartridge Space using Repack

Repack is an operation that recovers lost space on Data Cartridges by copying files from one Volume to another, omitting 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 cartridges. As well as recovering space that is wasted by old versions of files, it is also used to move data from one cartridge format to another (for example when a new, higher capacity format becomes available). The repack operation does 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.

- Files are copied to target destinations defined by the current [File Group rules](#). A File Group rule must exist for all files that are stored on the Volume that is being repacked.
- When all the files on the Volume have been successfully repacked, the repacked cartridges are moved to the [Quarantined Cartridge](#) 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 or optical library or at least two stand-alone drives. The repack operation cannot be performed on an archive with only one stand-alone tape or optical drive.

2.14 Alternate Data Streams

Alternate data streams, also known as 'NTFS streams' and 'named streams', are additional data streams that can be associated with a file. The relationship between a file and an alternate data stream is conceptually very similar to the relationship between a directory and a file; just as a directory can contain zero or more named files, a file can contain zero or more named alternate data streams. An alternate data stream is accessed by appending ':' (colon) and the stream name to the name of the file. 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.

Support for alternate data streams in XenData6 Server depends on the data cartridge format in use. The TAR (tape) cartridge format fully supports alternate data streams. The LTFS (tape) and ODA/UDF (optical) cartridge formats do not support alternate data streams because the underlying format specifications do not support alternate data streams. While it would technically be possible to support streams in the XenData6 Server implementation, attempting to do so has been found to break compatibility with other implementations of these file systems.

XenData6 Server from version 6.20 supports a range of commonly used alternate data streams by using [Hidden File Group Policies](#) to ensure that alternate data streams with common names are not written to the archive and do not cause errors with the LTFS and ODA file systems.

Alternate data streams 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 including the share of a volume managed by XenData6 Server. These alternate data streams contain application-specific metadata and/or Finder display layout information. XenData6 Server software employs special automatic file group rules that preserve Finder display information on the cache volume but do not write it to data cartridges. By default, XenData6 Server will attempt to write other types of application-specific alternate data stream to the archive and if this is not supported by the underlying cartridge format then these types of files cannot be written to the archive. Use of alternate data streams can be disabled on the Mac client, in which case the client writes 'AppleDouble' files consisting of a data fork file and a separate resource fork file. Writing files in this format is fully supported by all cartridge formats supported by XenData6 Server software.
- Windows Internet Explorer adds a stream named 'Zone-Identifier' to files downloaded from the Internet. Windows uses this data for security purposes. XenData6 Server software preserves this information on the cache volume but does not attempt to write it to data cartridges.

Alternate data streams are not visible in Windows Explorer, and their size is not included in the file's size. Alternate data streams may be listed on a Windows computer from the command prompt using the 'DIR' command with a '/R' parameter.

Example:

Typing the following command into a Windows command prompt will create a file with an alternate data stream...

```
echo Hello, world>File001:stream.txt
```

.. this command reads the alternate data stream...

```
more < File001:stream.txt
```

...and this command lists the alternate data streams in a file

```
dir /R
Volume in drive X is Archive.
Volume Serial Number is 48C9-8625

Directory of X:\XenData

03/07/2015  15:53                0 File001
                14 File001:stream.txt:$DATA
```

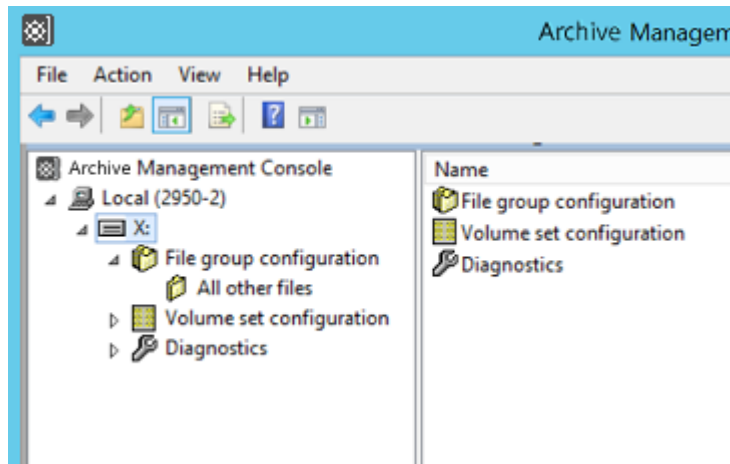
Note that the size of the file is shown as 0; in this case only the alternate data stream contains data.

Administering the System

3 Administering the System

The Archive Management Console is used to configure all [File Group](#), [Volume Set](#) and cartridge replication options, and to view diagnostic information about the [system](#) and [hardware](#). The Archive Management Console is a Microsoft Management Console (MMC) snap-in.

The Archive Management Console is shown below:



To start the Archive Management Console

If the operating system is Windows 7 or Windows Server 2008R2:

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

If the operating system is Windows Server 2012R2:

1. Click the Windows Start icon at the lower left hand corner of the screen.
2. Type "XenData System Configuration"
3. Click the **XenData System Configuration** entry in the list to the right of the screen.

To navigate around the Archive 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 under control by 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.

To determine the software version

1. Right-click on the drive letter in the left pane and then click System Details.

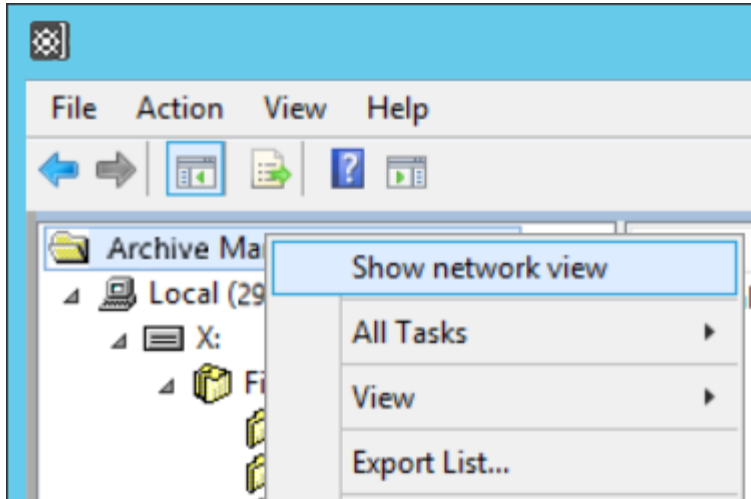
The version of the installed XenData6 Server software will be displayed.

To display Archive Management Consoles for other archives on the network

1. Right-click on the text Archive Management Console at the top of the left pane and enable Show

network view.

The left pane then displays the Archive 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 New Installation Check List

After installing XenData6 Server software on a new system, the Administrator must define all [File Groups](#) and [Volume Sets](#) before writing any files to the system. 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).
For each File Group, determine:
 - If files are to be saved to data cartridge Volumes or retained on the cache disk
 - For files saved to data cartridges, which Volume Set will be used, along with its associated configuration
 - If file fragmentation is to be enabled and, if so, the fragment size
 - For files saved to data cartridges, whether the files are to be flushed from the cache disk and if so, for what period should they be retained on disk after writing
 - For files scheduled to be flushed from the cache disk, whether they will be retained on disk after reading, and if so, for what period
- How many Volume Sets are required? (Remember that a Volume Set is a set of data 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 (optical or tape, WORM or rewritable) of data cartridges
 - For rewritable tape cartridges, the file system format, either TAR or LTFS
 - If the 'Write to disk if no writable volumes 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

Care should be taken to match the system settings to the capabilities of the storage hardware, especially the cache disk capacity, the use of a robotic library and/or stand-alone drives and the number of drives in each library. It is important to match the capacity of the cache disk and the file flushing policies, as it will not be possible to archive or restore files if the system runs out of cache disk space.

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

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

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

3.1.1 Special Considerations for Creating Directories at the Root of the Archive

The default file group rule that is created during initial system installation has the "Save files to archive" option selected and pointing to the default Volume Set. If no volumes have been added to the default Volume Set then it will not be possible to create a folder (other than "New Folder"). This is a consequence of XenData6 Server support for workflows that involve transporting data by physically moving cartridges from one location to another. Many such workflows require that empty folders are transported along with data files. Therefore the system preserves empty folders by writing them to data cartridge Volumes and if there are no writable Volumes available then an empty folder cannot be created.

Many system administrators wish to organize the archive by using different sub-folders of the root for different types of data (e.g. different customers etc.) By doing this and creating an appropriate set of File Groups and Volume Sets it is possible to specify that each section of the archive is stored on a different Volume Set. In this situation, the first thing that a system administrator will need to do is create a set of sub directories at the root of the archive, and this operation will fail because of the lack of writable cartridges.

The way to avoid this problem is to temporarily change the default file group rule so that the "Save files to archive" option is not enabled. The initial folder structure may then be created before changing the file group rules back to their final configuration.

3.2 File Groups

A File Group is a collection of files that all have the same file management policy and consequently are all treated in the same way by the system. Whenever a file is used, XenData6 Server software needs to know how to handle it. This 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 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 a combination of both. **Note** that a file's File Group is determined by the rules in place each time the file is used. It is not a persistent property of a file.

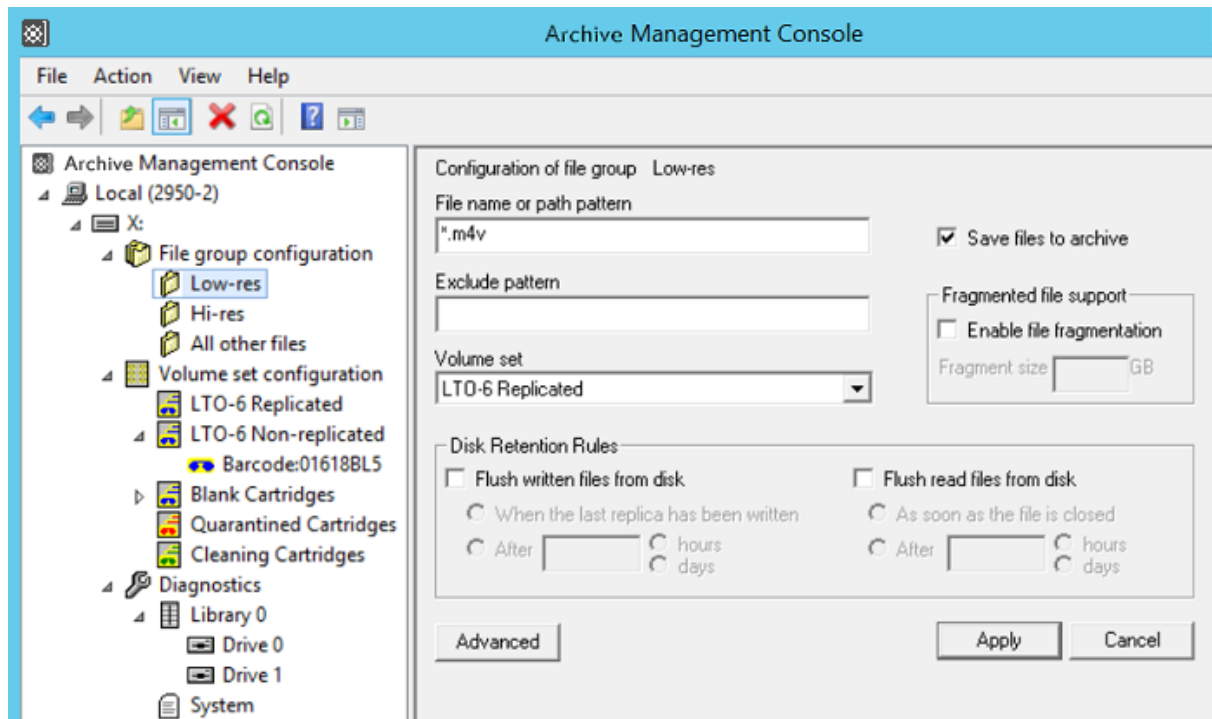
After initial installation of the XenData6 Server software, the system is configured with a single File Group called "Default". Typically the administrator will [Set Policies](#) for the Default File Group and perhaps [Create New File Groups](#).

3.2.1 Allocating Files to a File Group

To Allocate Files to a File Group

1. Open the Archive Management Console.
2. Navigate to the File group.
3. Update the File, name, or path pattern box with the appropriate text. For help on what to put in the box, click [here](#).
4. If required, update the Exclude pattern box.
5. Click **Apply**.

The order of File Groups in the left pane of the Archive Management Console is important and affects how files are allocated to File Groups (see [Changing the Order of File Groups](#)).

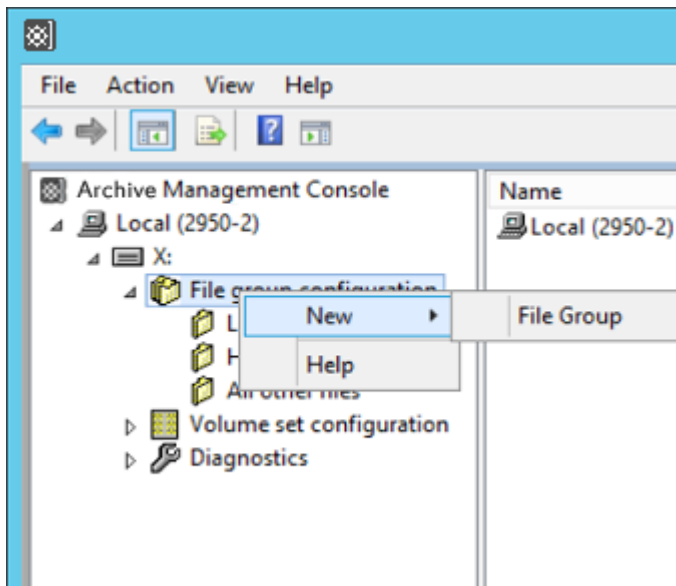


3.2.1.1 Creating a New File Group

To Create a New File Group

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

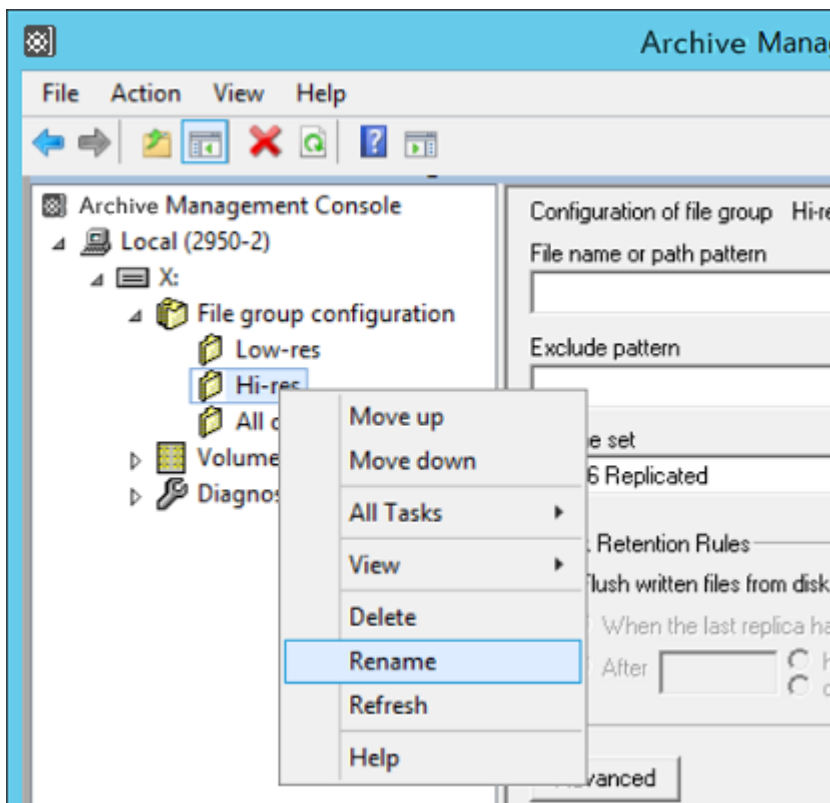
The new File Group will be named <<new>> and should be renamed as described in [Renaming a File Group](#). It should then be edited as described in [Allocating Files to a File Group](#) and [Selecting Storage Options for a File Group](#).



3.2.1.2 Renaming a File Group

To Rename a File Group

1. Open the Archive 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.

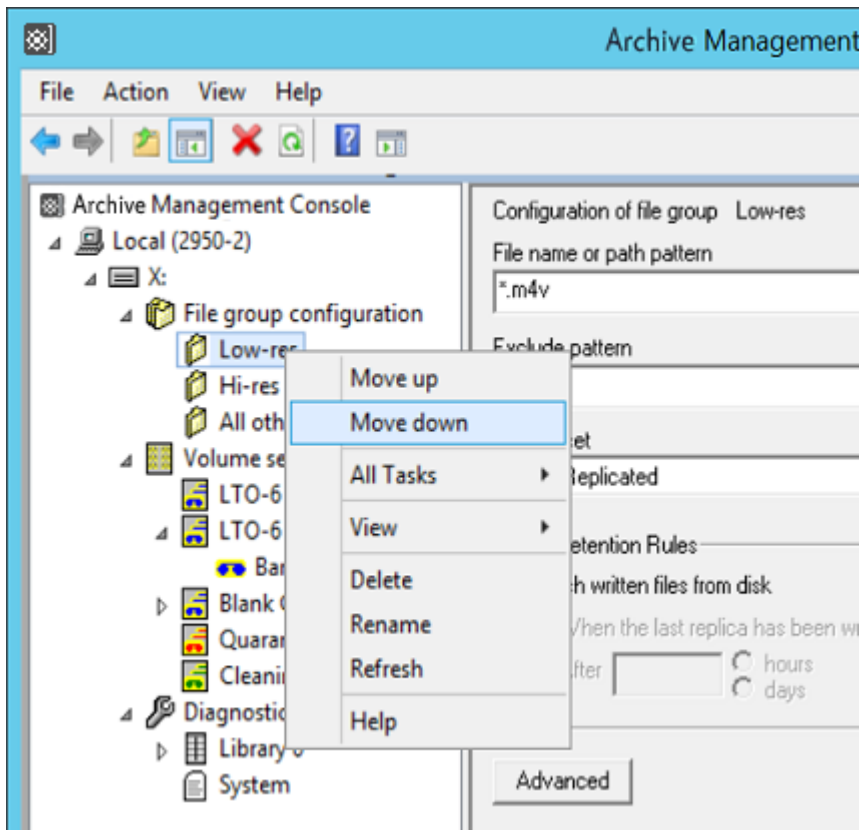


3.2.1.3 Changing the Order of File Groups

The order of File Groups in the Archive 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.

To Change the Position of a File Group in the List

1. Open the Archive 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**.



3.2.1.4 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:	*.tif
	Exclude pattern:	
File Group 2	File name or path pattern:	*.txt
	Exclude pattern:	
File Group 3	File name or path pattern:	*
	Exclude pattern:	

In this example, no path has been specified and consequently the file name rules apply to all files written to the archive, no matter which folder is used.

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 called “\project02\”; and a third File Group contains all other files.

File Group 1	File name or path pattern:	\project01*
	Exclude pattern:	
File Group 2	File name or path pattern:	\project02*
	Exclude pattern:	
File Group 3	File name or path pattern:	*
	Exclude pattern:	

Example 3: This is similar to example 2, but additionally includes all sub-folders of project01 and project02. One File Group contains all files written to “\project01\” and its sub-folders; another contains all files written to “\project02\” and its sub-folders; and a third File Group contains all other files.

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

Note that In this example, the use of the special pattern “...\” 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”. Consequently, all file names ending in “.tmp” are allocated to the third File Group.

File Group 1	File name or path pattern:	\project01\...*
	Exclude pattern:	*.tmp
File Group 2	File name or path pattern:	\project02\...*
	Exclude pattern:	*.tmp
File Group 3	File name or path pattern:	*
	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. As for example 4, all file names ending in “.tmp” are allocated to the third File Group. However, the administrator has not configured a ‘catch-all’ File Group rule at the bottom of the File Group list. In this example the system will not allow writing of files unless **either** they are written to the folders project01\, project02\ or their sub-folders **or** the file extension is “.tmp”.

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

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

File Group 1	File name or path pattern:	\project01*
	Exclude pattern:	
File Group 2	File name or path pattern:	*.tmp
	Exclude pattern:	
File Group 3	File name or path pattern:	*
	Exclude pattern:	

In this example, files with a ".tmp" extension 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 ".tmp" would be allocated to the same File Group as the ".tmp" files in the other folders.

3.2.2 Setting Policies for File Groups

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

- [Selecting Storage Options](#), including whether files in the File Group are saved to data cartridges in a Volume Set.
- Settings for [File Fragmentation](#).
- [Disk Retention Rules](#).
- Optional [Advanced Options](#).

3.2.2.1 How to Construct the "File Name or Path Pattern" or the "Exclude pattern" for a File Group

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 use the file).

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

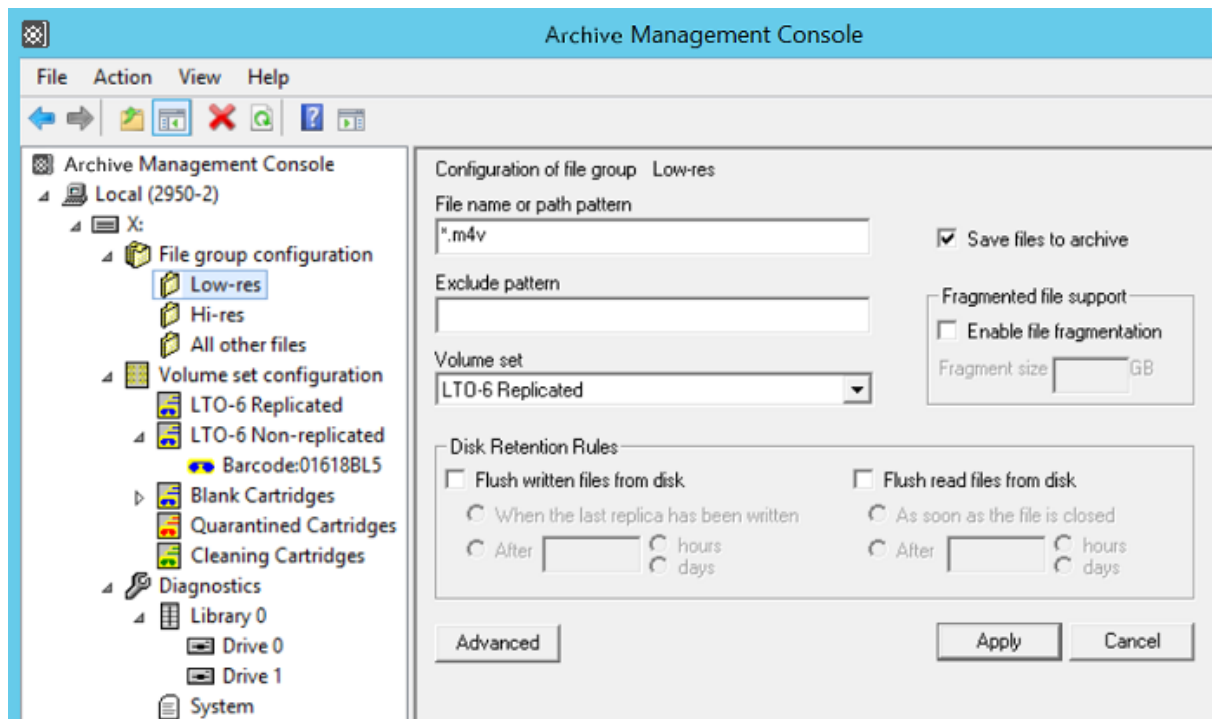
Some example file name or path patterns are:

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

3.2.2.2 Selecting Storage Options for a File Group

To Select Storage Options for a File Group

1. Open the Archive Management Console.
2. Navigate to the File group.
3. Determine whether files in the File Group are to be saved to data cartridges in a Volume Set. If so, enable the "Save files to archive" option and select the required Volume Set.
4. Determine whether or not file fragmentation is required for the File Group and enable if appropriate. File fragmentation is typically required only for large files, and is described [below](#).
5. If files are being saved to a Volume Set, determine whether to use file flushing to save space on the cache disk. If file flushing is to be used, configure the Disk Retention Rules as described [below](#).
6. If appropriate, click on the **Advanced** button near the bottom of the screen and make additional selections.
7. Click **Apply**



3.2.2.3 File Fragmentation

The term 'file fragmentation' refers to the way in which computer systems break large files into smaller, more manageable units for transfer to or from storage devices. Fragmentation is required with magnetic disks, because gaps are created when files are deleted. Fragmentation on a magnetic disk leads to performance degradation over time that can be corrected using de-fragmentation utilities. The tape and optical disc formats that are supported by XenData6 Server software do not inherently need to use file fragmentation because they use cartridge formats that are recorded sequentially from the beginning of the cartridge to the end, with individual files recorded as complete entities. However, optional file fragmentation is available and provides the ability to partially restore files, which can lead to significant performance improvements with large files. Furthermore, file fragmentation allows extremely large files to be spanned across multiple cartridge Volumes.

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 the archive, only the applicable fragments will be restored, saving both transfer time and space on the cache disk;
- when an application modifies a large file by appending, the appended data will be written to the archive as one or more additional fragments, saving space on the data 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 the archive, and
- on writing a file, the data may span multiple Volumes if the File Group advanced option to permit file spanning is enabled. 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 the archive, the whole file will be restored;
- when a file is modified, the new version of the file will be completely written to the archive
- on writing a file, the system will always write the whole 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 drive transfer rates but for LTO-7 will typically be 5 GB or larger.

Note: File fragmentation is not supported on ODA cartridges because it would break compatibility with the Sony tools and utilities.

3.2.2.4 Disk Retention Rules

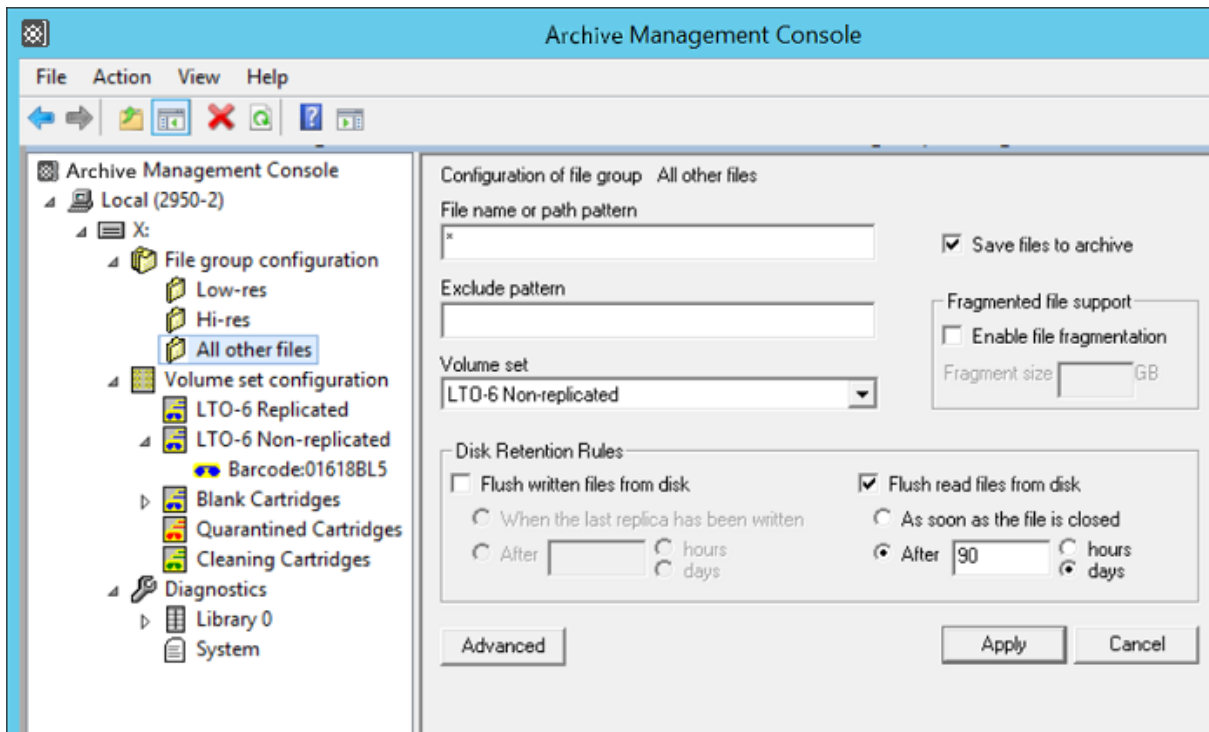
The administrator can configure the system such that after a file has been securely written to one or more data cartridges, the instance stored on the disk cache will be flushed to release the cache space occupied by the file. The flush operation causes the file to be removed from the cache, but it is still visible in the file system and accessible to normal applications from the archive. Flush 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 the "Save files to archive" option has been enabled.

Characteristics of flushed files are as follows:

- Flushing from the cache does not affect the presence and location of a file within the file system.
- File properties - including file size, modification date etc. - do not change, except that the Windows offline attribute bit is set.
- Flushed files are restored by simply reading the file. As long as the file is available on a near-line data cartridge, the system will restore the file automatically.

To Configure Disk Retention Rules

1. Open the Archive Management Console.
2. Navigate to the File group.
3. Enable the "Save files to archive" option.
4. Update the Disk Retention Rules, as described below.
5. Click Apply.



When a file is first written to the XenData6 Server system, it is first written to the cache disk. If all the options in the Disk Retention Rules area are left unselected, the system will retain the most recent versions of all files on disk indefinitely. The options that are described below allow the administrator to flush files from the cache disk when the file has been written to one or more data cartridges. **Note** that the system takes great care not to flush files until the last replica has been securely written to the archive.

If a file that has been flushed from the cache disk is subsequently read, the system restores the file to the cache disk. The administrator can use additional options described below to flush restored files from disk when the user has finished with them.

Options for Disk Retention

- Files are retained on disk indefinitely. Deselect both options in the Disk Retention Rules. This option is used for small, frequently accessed files such as certain types of application specific metadata.
- 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 securely written to the archive. However, unless the **Flush read files from disk** option is also selected, any files that are subsequently restored from the archive will be retained on disk indefinitely.
- Files are flushed a preset length of time after writing. Select **Flush written files from disk** and **After** and choose a number of hours or days. With this option selected, files are retained on disk for the defined length of time after they were written, or until the last replica has been completely written, whichever comes later. As above, the **Flush read files from disk** option is also usually selected to prevent files that are subsequently read back from the archive from being retained indefinitely.
- Files are flushed immediately after being read. Select **Flush read files from disk** and **As soon as the file is closed**. This option is generally used in combination with the **Flush written files from disk option**. In that case files will be flushed from disk immediately after they have been read unless the period specified in the **Flush written files from disk** option has not yet passed, in which case they will be retained for the remainder of the time specified by that option. If the **Flush**

read files from disk option is selected without **Flush written files from disk** option then files will be retained on disk until the first time they are read, however long that takes. This combination of options is sometimes useful where some sort of post-processing is performed on files after they have been written to the archive.

- e. Files are flushed a preset length of time after being read. Select **Flush read files from disk** and **After** and choose a number of hours or days. As in d) above, this option is generally used in combination with the **Flush written files from disk** option and the interactions between the two options are similar to the case above; files are retained on disk for the longest of the two time periods specified (read flush interval and write flush interval). **Note** that the **Flush read files from disk** timer is reset every time a file is read, so if files are read frequently, the effect of this option can be that they are never flushed.

3.2.2.5 Changing Retention Rules

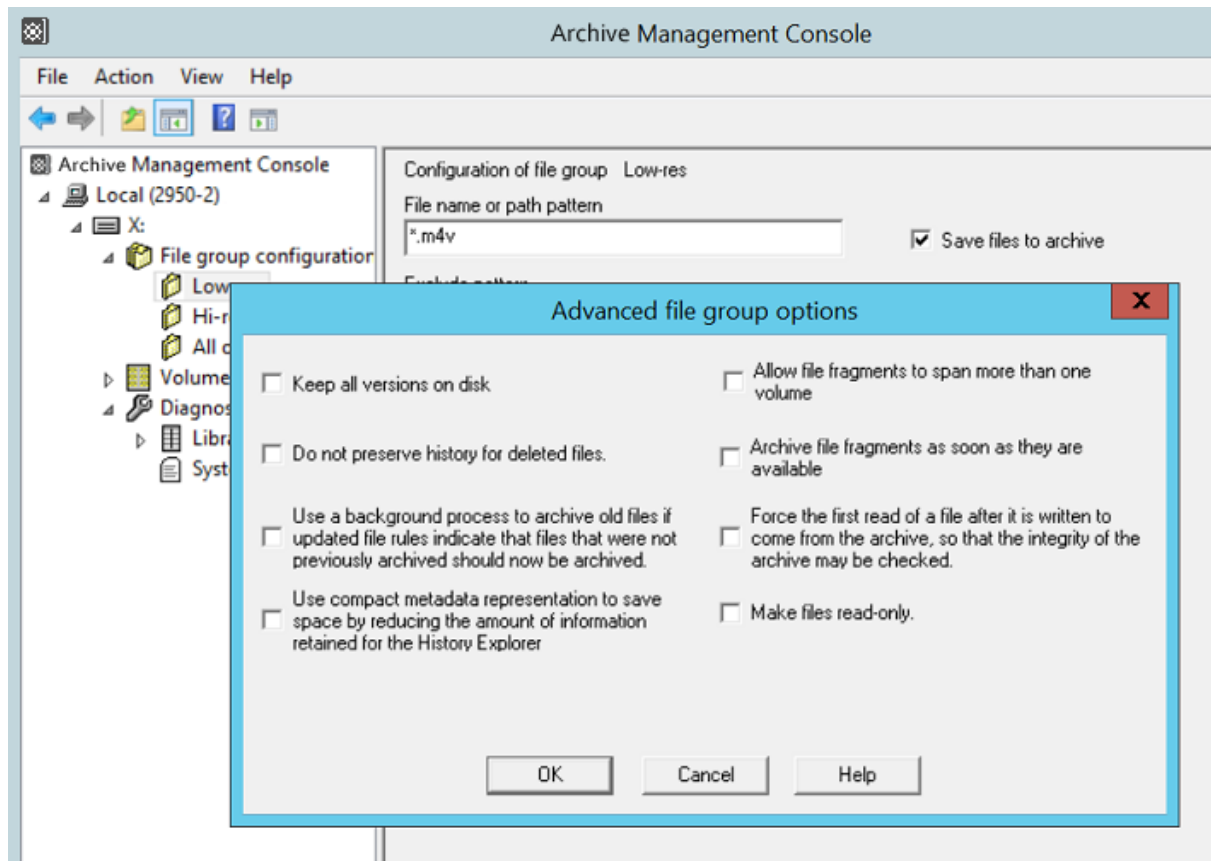
The system administrator may change Disk Retention Rules for a File Group at any time during system operation. If the rules are changed, the new rules apply to all files in the File Group, not just to new files that are created after the rule change is implemented. Thus, if a system is running short of space on the cache volume, the system administrator can change retention rules to keep files for a shorter length of time and the system will immediately start to free space by flushing old files. (Alternatively, files may be flushed using the Windows Explorer Flush operation).

3.2.2.6 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.

To Configure File Group Advanced Options

1. Open the Archive Management Console.
2. Navigate to the File group.
3. Click on the Advanced button near the bottom of the screen.



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 History Explorer. To conserve space, the system removes old versions of files from the disk cache, making them available only from data cartridges. Selecting the "Keep all versions on disk" option changes this behavior so that all old versions of a file are retained on the cache.

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. This functionality is required for many applications including compliance and anywhere there may be a need to recover data that has been accidentally deleted or overwritten. However, a limited range of applications produce large numbers of intermediate files that are valuable at the time they are created, but at some later time 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 archive (for example, some email archiving systems work this way). Because the incoming files (for example, individual email messages) are valuable, they should be archived, providing an immediate backup. 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 cache disk (for metadata) and if the system is maintaining metadata for a very large number of deleted files, the space consumed 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.

A second use of this option relates to situations where it is necessary to delete files from the cache

volume without deleting them from data cartridges, perhaps because the data cartridges have been permanently moved to a different location. In this case, selecting the option will allow files to be deleted without attempting to access the cartridges that contain the data.

Note: 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
In normal operation, XenData6 Server software writes a file to the Volume Primary replica as soon as it has been closed. By doing this, the system provides the improved data integrity and security necessary for an archival data storage system by ensuring that a second instance of all data written to the system is generated almost immediately. In the case of WORM media, it further ensures that an unalterable copy of the data is made in a timely fashion, as required by many regulatory authorities. However, in the event of a shortage of blank cartridges or a hardware failure, no data can be written to the system. For high-integrity compliance applications that require data to be written immediately to WORM media there is no alternative but to stop writing data until the failure condition has been resolved. In other applications it may be acceptable for a short time to write data only to the disk cache, waiting until the fault has been resolved before writing the data to archival data cartridges. A system administrator can achieve this by turning off the "Save files to archive" file group option. When the failure has been resolved, the system administrator re-enables the "Save files to archive" option and enables the "use a background process..." option as well. New files will be written to data cartridges in the usual way and at the same time the system performs a background scan searching for older files that should now be archived.
- 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 History Explorer is the creation date for every version of every file. This information is stored in the file's metadata on the cache disk and it is only used by History Explorer (or other applications that use the API). If a very large number of versions of a file are created, this extra information increases the size of the metadata, consuming space on the cache volume. This option changes the format of the metadata so that file version modification dates are not saved, thereby reducing the size of the metadata. This optimization is not worthwhile for most archiving applications because the number of file versions created is usually small, the metadata file size is smaller than the cache disk cluster size and no space will actually be saved on the cache volume.
- e) Allow file fragments to span more than one volume
This option is applicable when file fragmentation is enabled. It determines whether or not an individual file's fragments may be written so as to span across multiple data cartridge Volumes. When this option is not enabled, all file fragments for a particular version of a file will be written to the same Volume.
- f) Archive file fragments as soon as they are available
In normal operation, XenData6 Server software writes files to the archive after the whole file has been written to the cache disk and the file has been closed. 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 is sometimes more efficient to write data to data cartridges as soon as the application has finished writing each 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 default behavior of the system is always to read data from the fastest available location. For data that has just been written to the archive, this will usually be the cache

disk (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 data cartridges. This option forces data to be read from data cartridges even if it is available from an intermediate cache, thereby allowing applications to verify the integrity of data on the cartridges in the archive. **Note** that only the first read is forced to come from the cartridges; subsequent reads will be satisfied from the cache if possible.

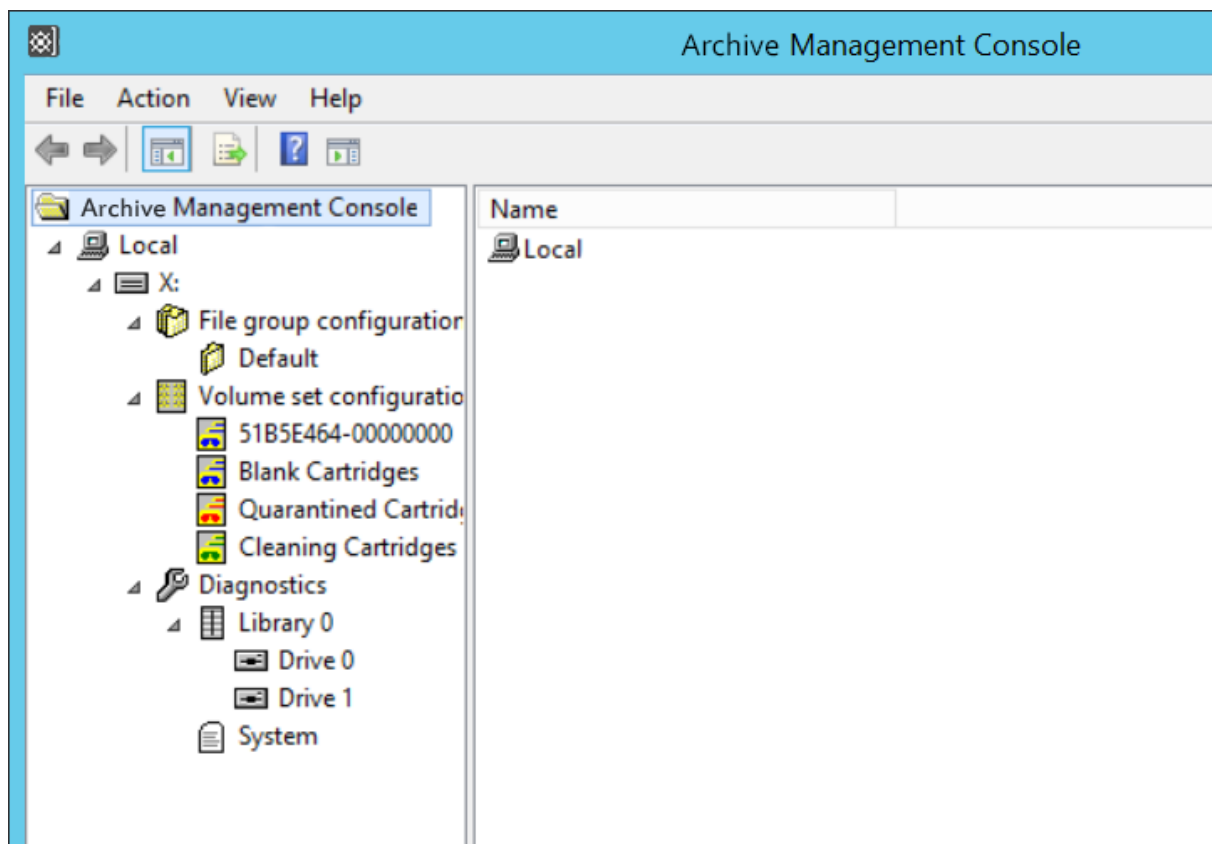
h) Make files read-only

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

3.3 Volume Sets

A Volume Set consists of a set of data cartridges with a defined cartridge type (tape or optical disc cartridge, [WORM](#) or [rewritable](#)) and file system format ([LTFS](#), [TAR](#) or [ODA](#)). Each Volume Set stores files from one or more [File Groups](#). Automatic cartridge [replication](#) can be defined for tape based Volume Sets.

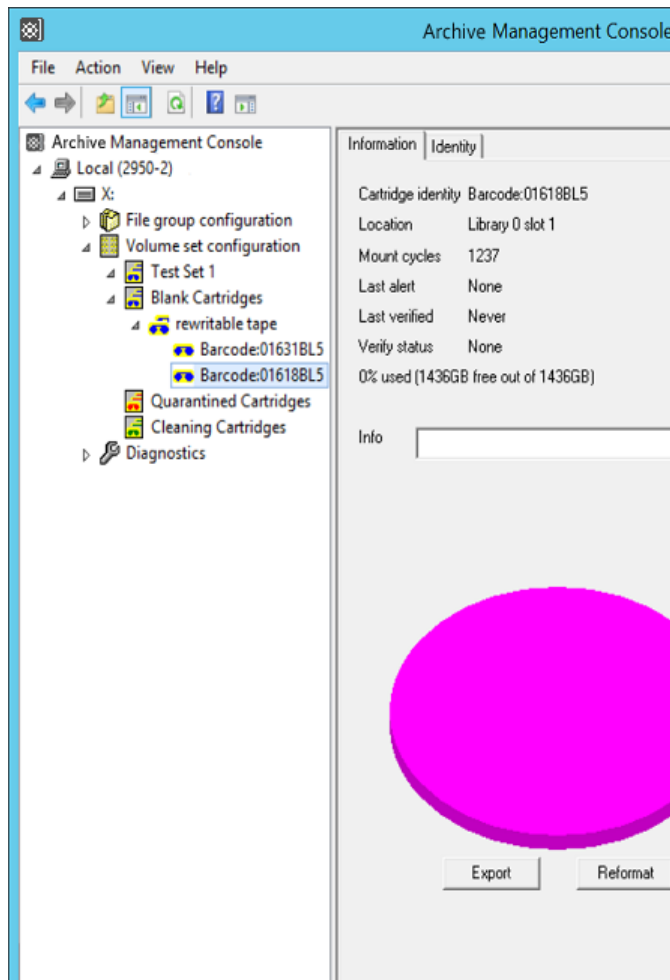
When a new system is installed, an initial Volume Set is created, ready for [configuration](#). In addition, three special Volume Sets are shown in the Archive Management Console: the [Blank Cartridge Set](#), the [Quarantined Cartridge Set](#) and the [Cleaning Cartridge Set](#).



3.3.1 The Blank Cartridge Set

The Blank Cartridge Set is a special [Volume Set](#) that contains data cartridges that have been imported into a robotic library or inserted into a stand-alone drive but are not yet 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 Cartridge Set are allocated to an operational Volume Set either manually by the administrator or automatically by the XenData6 Server software when all the existing Volumes in a Volume Set are nearly full. The threshold at which new Volumes are added to a Volume Set is determined by the Volume Set configuration settings.

Volume Set configuration settings also determine the behavior of the system when all existing Volumes in a Volume Set become full and no more blank cartridges of the correct type are available. 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 write data to the cache disk but not to data cartridge Volumes. 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.

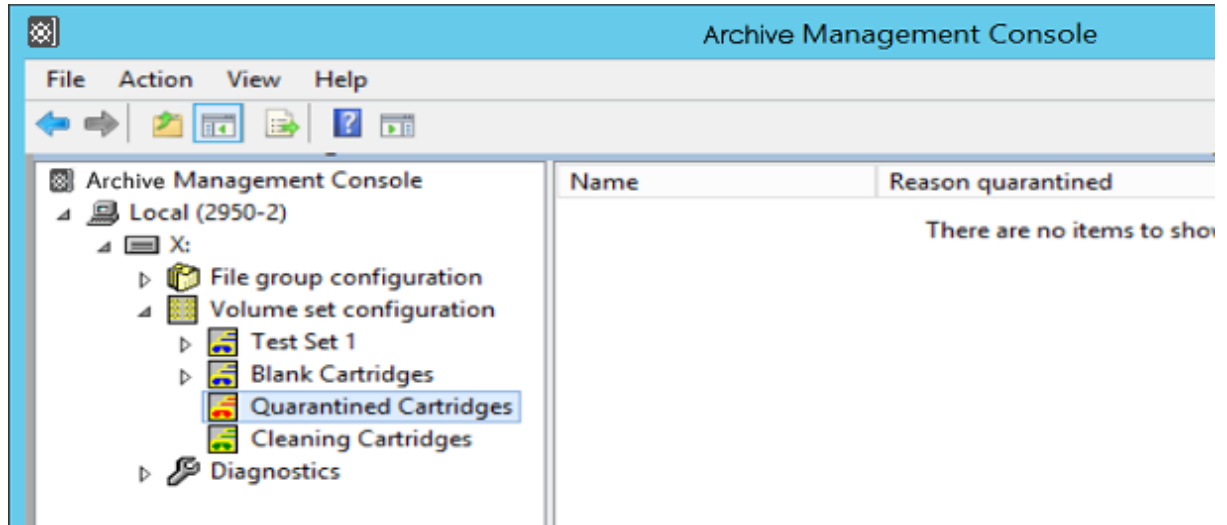


3.3.2 The Quarantined Cartridge Set

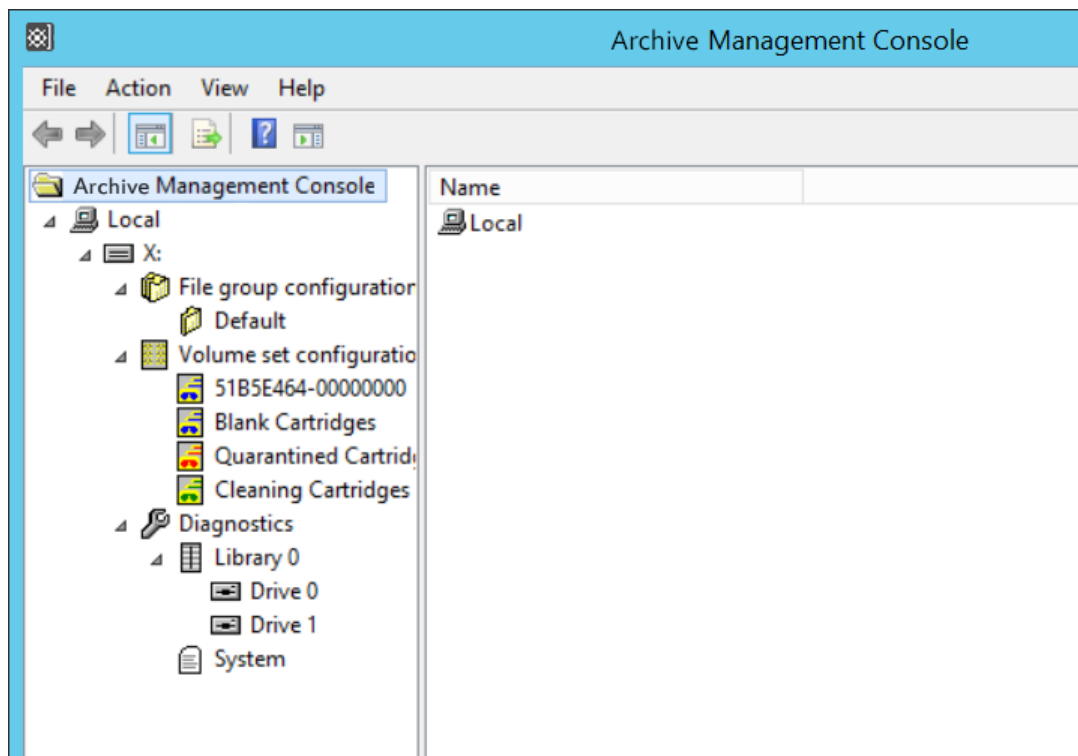
The Quarantined Cartridge Set is a special [Volume Set](#) that contains data cartridges that have been imported into a robotic library or inserted into a drive but for some reason cannot currently be used by

the system. Typically, this will be because the cartridge contents have been repacked, because the cartridge has previously been formatted by an incompatible application (such as a backup application) or because an error occurred while the system was trying to identify the contents of the cartridge.

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



3.3.3 Volume Set Defaults for a New Installation



After a new installation without any data cartridges in the library or stand-alone drive, the Archive Management Console will look as shown above.

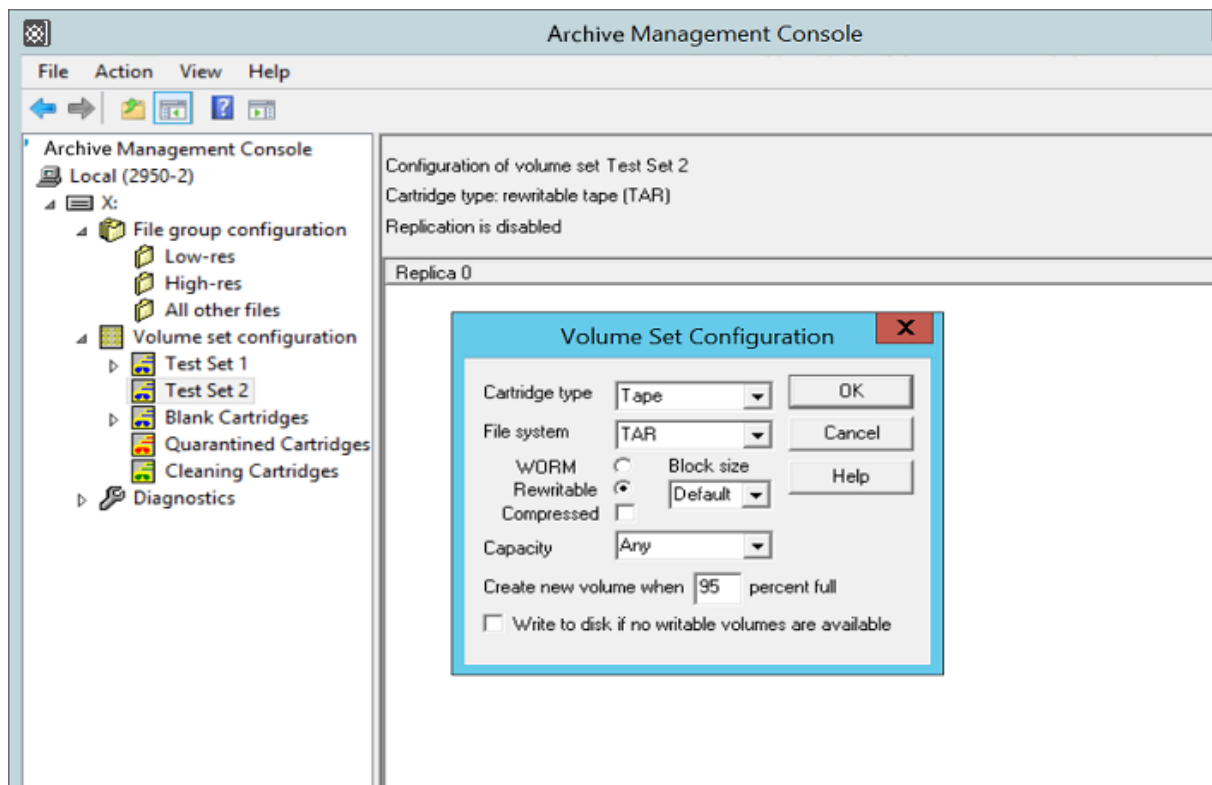
With a new installation, the default system will have an initial default Volume Set, a blank cartridge set, a quarantined cartridge set and cleaning cartridge set. The default Volume Set will have a

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 configured by selecting it in the left pane and then using the 'configure' and 'replication' buttons in the right pane, as described below.

3.3.4 Configuring a Volume Set

To Configure a Volume Set

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



Basic Volume Set configuration

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

1. Select **Tape** or **Optical** cartridge type, as appropriate for your hardware.
2. Select **WORM** or **Rewritable** data cartridges using the two radio buttons.
3. For tape cartridges, select the format, either **TAR** or **LTFS**.
4. The **Compressed** check box enables data compression if the hardware supports it. If the system is storing uncompressed data then selecting this option will allow an increased amount of data to be saved per cartridge. However, many applications perform their own data compression and if this is the case then it is unlikely that the hardware compression built into the drive will offer any further compression and may increase the file size because of compression overhead.
5. Select the **block size** for Volume Sets that use the TAR (tape) 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 hardware.
6. You can also define a cartridge **capacity** and change the point when additional Volumes will be

automatically **created** from the [Blank Cartridge Set](#).

7. The **Write to disk if no writable volumes are available** option determines the system's behavior if all Volumes in the Volume Set become full or unavailable. If this option has been enabled and all Volumes in a Volume Set become full or unavailable, the system automatically enters the [Pending Write Mode](#) and will accept more data which is stored on the cache disk. 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.

After Configuring the Volume Set

If required, [configure replication for the volume set](#).

When you are satisfied with the Volume Set configuration, add media to the Volume Set, as described [here](#). The Volume Set is then ready for use.

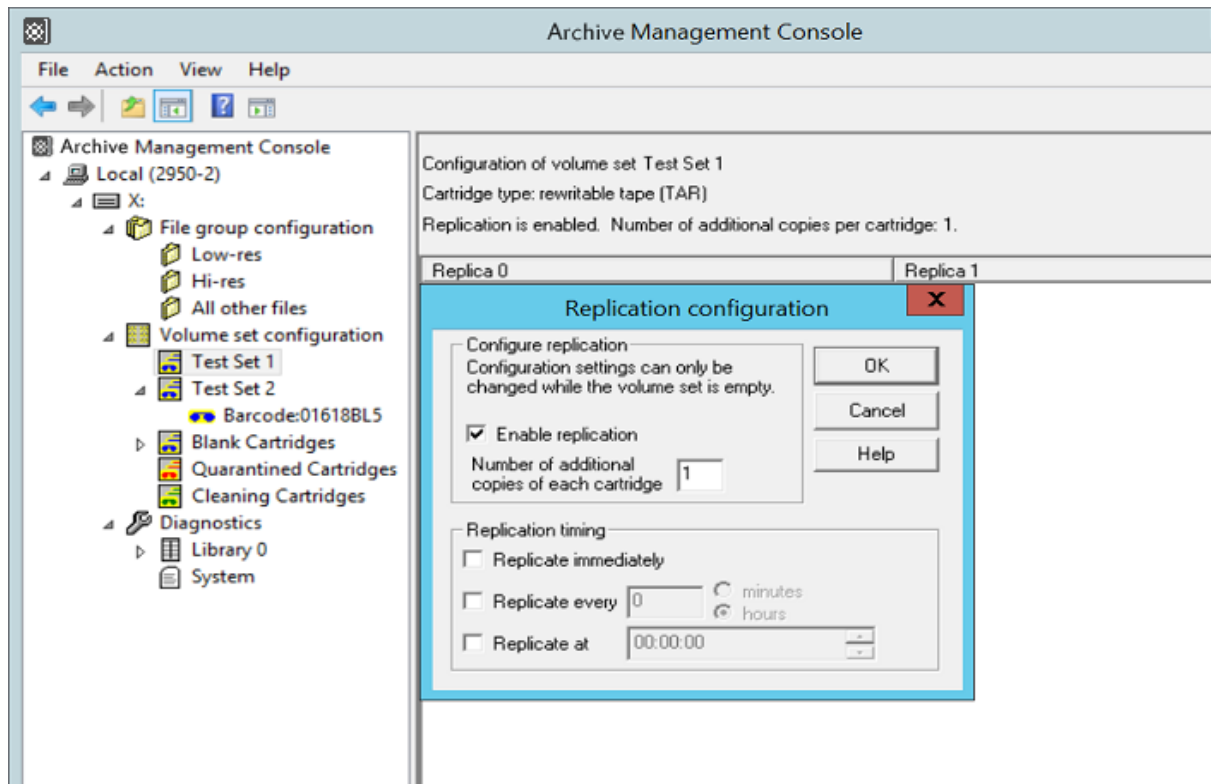
Note that once a Volume is added to a Volume Set, the Volume Set cannot be deleted and the data cartridge type and number of replicas cannot be changed.

3.3.5 Configuring Replication for a Volume Set

Replication settings for a Volume Set must be defined before manually adding the first Volume. After adding a Volume, only the replication schedule can be changed, not the number of replicas. **Note** that Volume Set replication is not supported for ODA optical disk cartridges or on a system that has no library and only a single stand-alone tape drive.

Replication is defined for a Volume Set as follows:

1. Open the **Archive Management Console**.
2. Navigate to the **Volume set configuration** section.
3. Select 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 box, as shown below.
5. Check the **Enable replication** box.
6. Enter the required number of replicas in the **Number of additional copies of each cartridge** box.
7. Then define [Replication Timing](#). Replication timing schedules can specify that replication should take place immediately; periodically or at the same time each day.
8. Click on **OK**.



3.3.6 Defining 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 Replication Timing:

1. Open the Archive 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 box.
5. Define the required replication schedule.

To define an "immediate" replication schedule

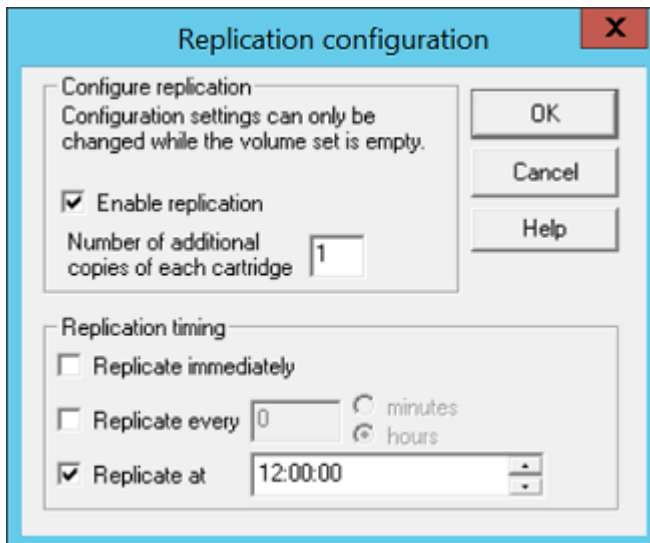
Check the Replicate immediately box. Note that immediate replication cannot be selected in conjunction with any other replication schedule. Immediate replication is not recommended for single drive library systems, because it will cause increased cartridge swapping, leading to degraded performance.

To define a "periodic" replication schedule

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

Enter the required replication time using the time picker. The system supports simultaneous selection of the "Replicate every" and "Replicate at" options.

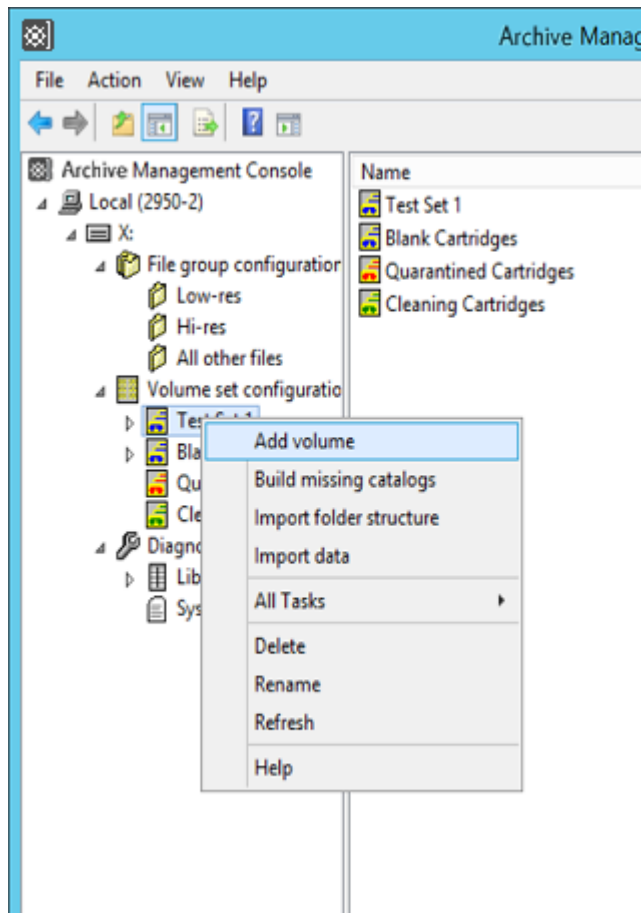


3.3.7 Allocating Volumes to a Volume Set

After initial creation and configuration of a new Volume Set, the system administrator must add cartridges as described below. In a system that uses a robotic library, blank cartridges are automatically added to the Volume Set when all the volumes assigned to a particular Volume Set are nearly full (as defined in [To Configure a Volume Set](#)).

To Add Volumes to a Volume Set

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

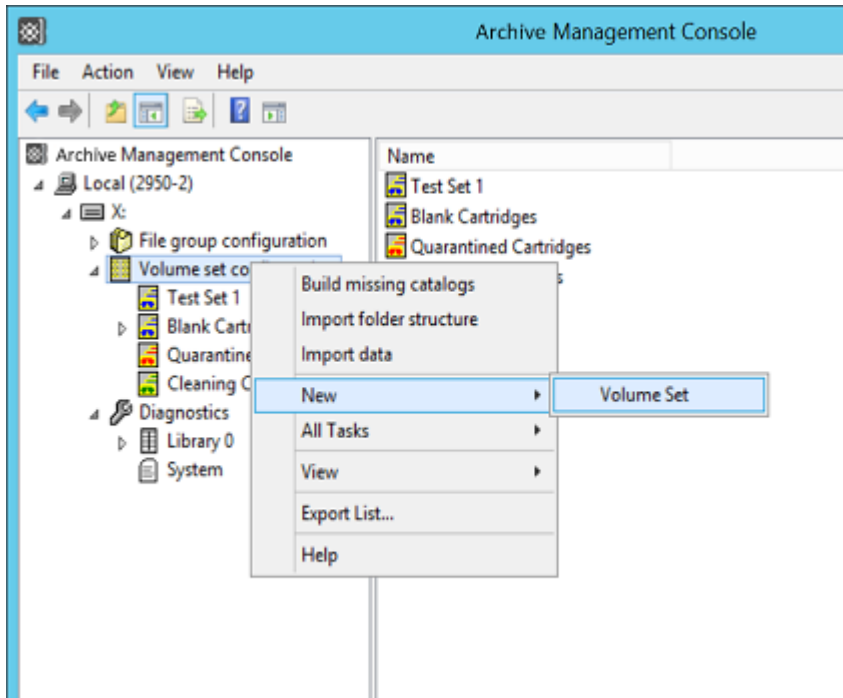


Recommendation: Do not preallocate more than one Volume to a Volume Set. It is better to leave the system to automatically allocate Volumes as required. Preallocating additional Volumes unnecessarily increases the time taken for the system to initialize.

3.3.8 Creating a New Volume Set

To create a new Volume Set:

1. Open the Archive 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 will be 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.

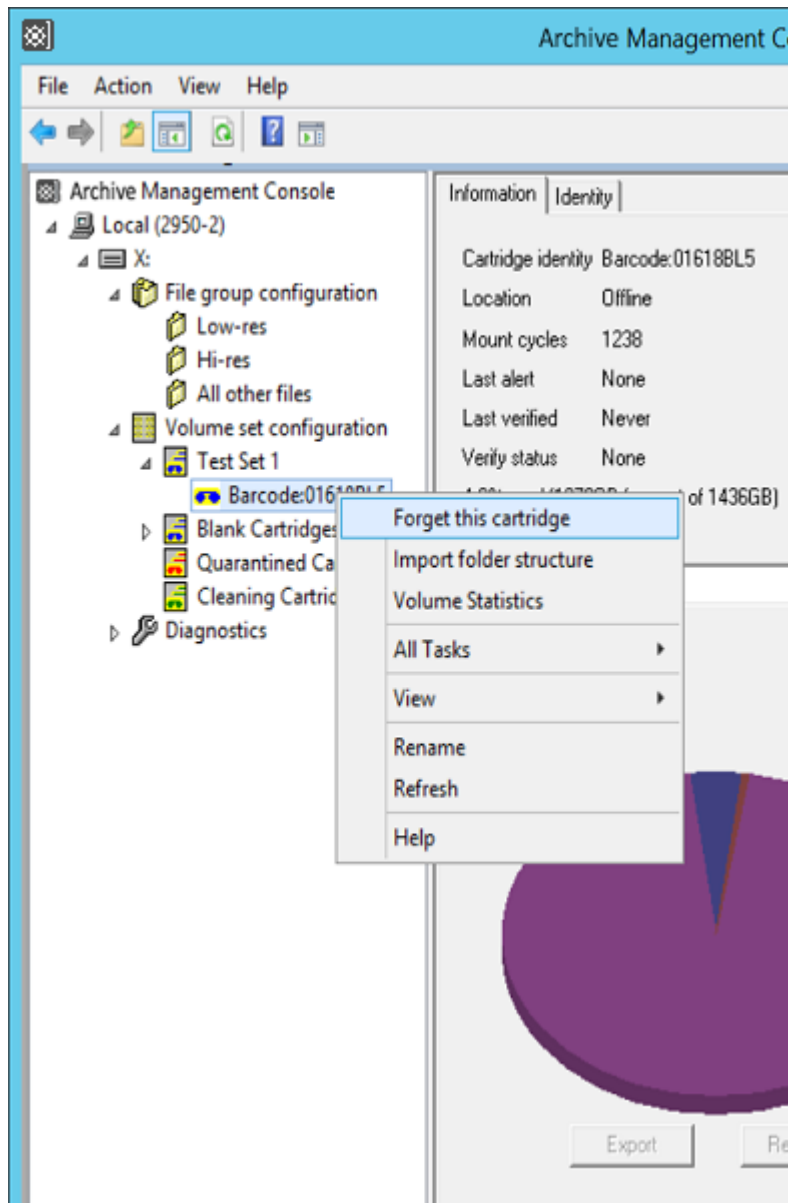
The new Volume Set should then be configured as described in [Configuring a Volume Set](#), [Defining Replication for a Volume Set](#) and [Allocating Media to a Volume Set](#).

3.3.9 Removing Information about a Cartridge from the System

If a replica cartridge becomes lost or damaged or if it is permanently moved to a different location, it might be convenient to remove information relating to that cartridge from the system. It is necessary to do this when replacing a damaged data cartridge in a replica set. This is called "Forgetting" a cartridge.

To Forget a Cartridge

1. Open the **Archive Management Console**.
2. Navigate to the cartridge.
3. Ensure that the cartridge is offline.
4. Right click on the cartridge and select **Forget this cartridge**.

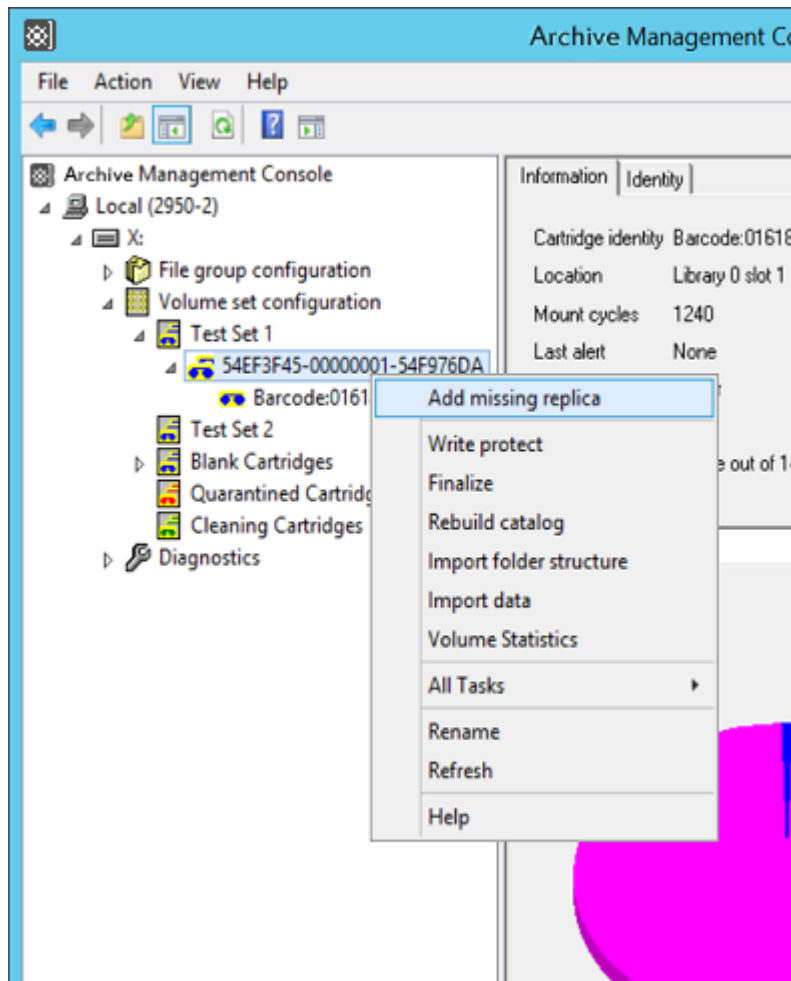


3.3.10 Replacing a Missing Replica Cartridge

If [Replication](#) is enabled for a [Volume Set](#) then files are automatically archived to two or more replica cartridges. If a replica cartridge becomes lost or damaged, it can be replaced using the **Add missing replica** operation. However, before this operation can be used, the system must be instructed to [Forget](#) the missing replica.

To Add a Missing Replica

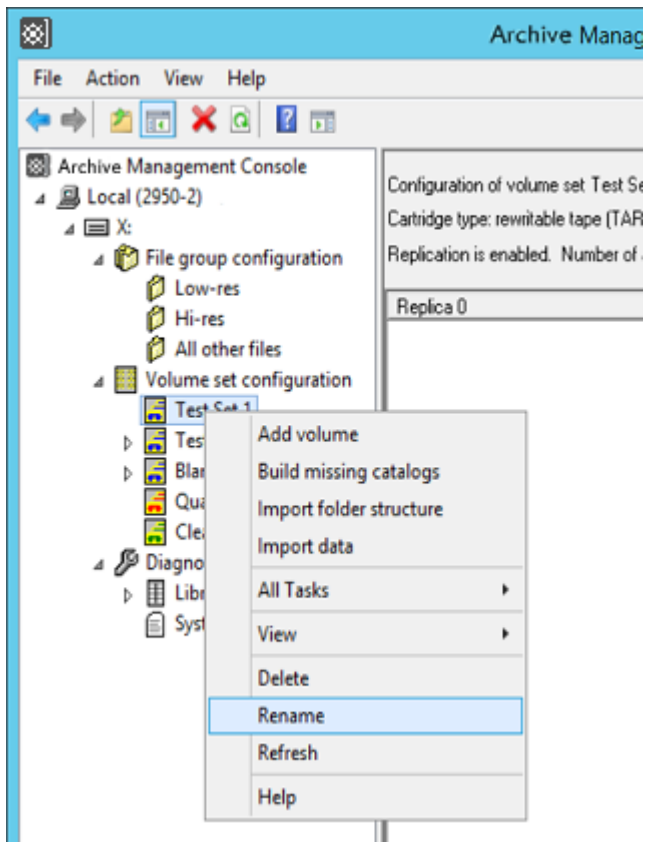
1. Open the Archive Management Console.
2. Navigate to the Volume.
3. Right click on the Volume and select Add missing replica.



3.3.11 Renaming a Volume Set

To Rename a Volume Set

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

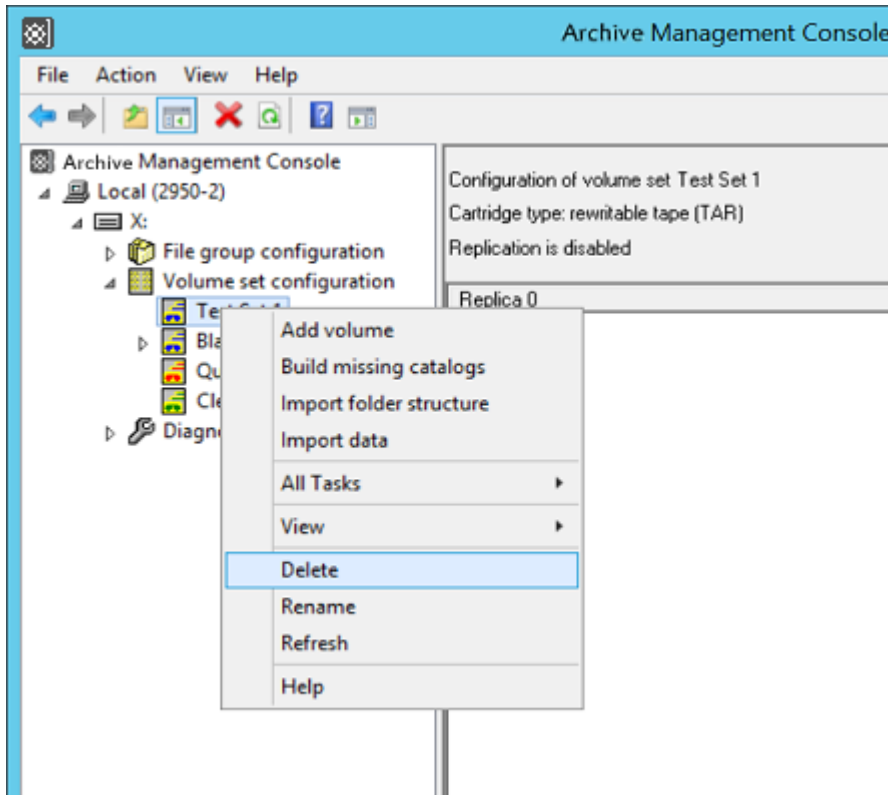


3.3.12 Deleting a Volume Set

To delete a Volume Set:

1. Open the Archive 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 Volumes and it has not been selected as the "Save to archive" target of any File Group.



3.4 Managing Data Cartridges

XenData6 Server software supports a wide range of stand-alone tape and optical disc drives and robotic libraries and can handle both [rewritable](#) and [WORM](#) cartridges. Please refer to your supplier for a current list of supported hardware.

3.4.1 About Rewritable Cartridges

Unlike [WORM](#) cartridges, data on rewritable cartridges can be overwritten or erased. This version of XenData6 Server software supports a range of rewritable data tape and optical disc cartridges. The software is licensed to use specific types and capacities of cartridges.

3.4.2 About WORM Cartridges

WORM is an acronym for Write Once Read Many. When data has been written to a WORM data cartridge, it cannot be erased or altered (additional data can, however, be appended to the cartridge). XenData6 Server software manages this non-rewritable medium in such a way that files can be 'deleted' within the file system in the normal way. In reality, deleted files are only being hidden from the normal file system interface and can be retrieved using the [History Explorer](#) utility which can 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. This version of XenData6 Server software supports a range of WORM tape and WORM optical disc cartridges. The software is licensed to use specific types and capacities of cartridges.

Note that the LTFS file system does not support WORM tapes because the LTFS format specification does not include support for WORM cartridges.

3.4.3 Importing Cartridges

Data cartridges are imported into a robotic library using the library controls, which vary from one model to another. In the case of a stand-alone drive, cartridges can be manually inserted into the drive. XenData6 Server software will automatically attempt to identify any new cartridge that is imported into the system.

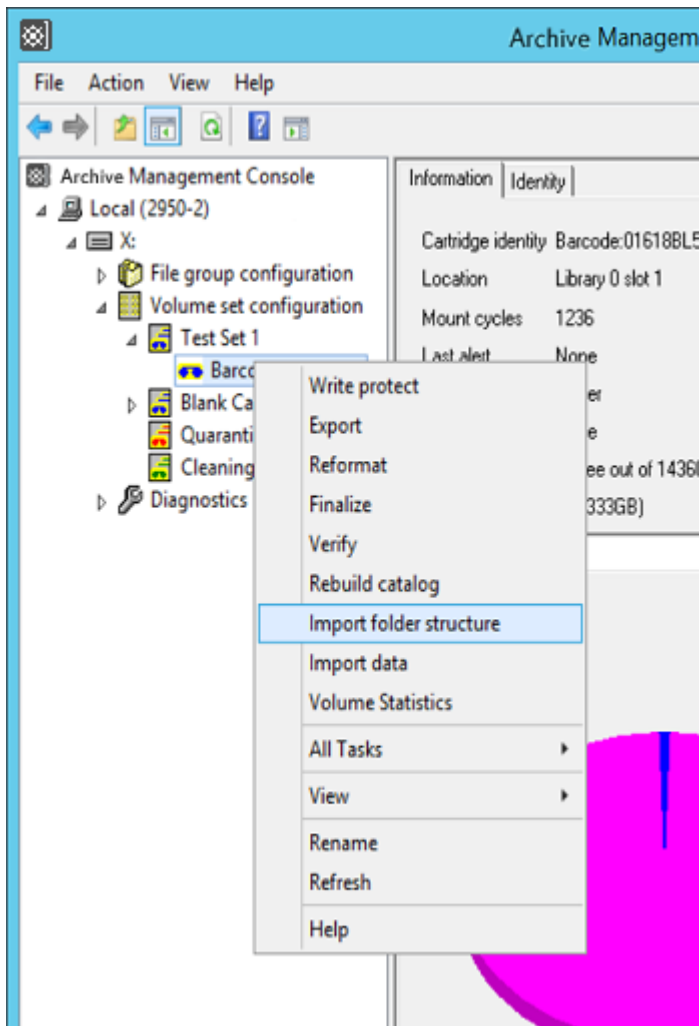
Note that if an unknown cartridge that contains data is imported into the system, files will not be visible until the Administrator uses either the **Import Folder Structure** or **Import Data** functions as described in the next section.

3.4.4 Transferring Cartridges Between Systems

When a new (previously unknown) cartridge is imported into a XenData6 Server system, the system examines the cartridge contents. If the cartridge has previously been written by XenData6 Server software it will be assigned to its original [Volume Set](#). If the original Volume Set does not exist, it will be created. If the cartridge is blank, it will be assigned to the [Blank Cartridge Set](#) and if the cartridge was written by a third party system in a format that is compatible with XenData6 Server software (ODA or LTFS) then it will be assigned to a special volume set with an identity that starts with FFFFFFFF. Cartridges written in any other format (i.e. cartridges written by a backup or other application that is not compatible with XenData6 Server software) will be assigned to the [Quarantined Cartridge Set](#).

Note: when a data cartridge is moved from one XenData6 Server system to another, the files on the cartridge will not immediately be visible in the file system interface. The administrator must use the **Import Folder Structure** function to make the contents of newly imported data cartridges available in the server file system. Files are imported from cartridges written on another system as described below:

1. Import the data cartridges into the library or insert a cartridge into a stand-alone drive.
 - The system automatically creates the Volume Sets and Volumes containing the newly imported cartridges.
2. Open the **Archive Management Console**.
3. Navigate in turn to each newly imported cartridge.
4. Right click on the cartridge and select **Import Folder Structure** or **Import Data**.
 - Import Folder Structure loads file and folder information (metadata) into the system, making the entire folder structure visible to users, but it does not restore any file instances to the disk cache. This operation is much faster than Import Data.
 - Import Data also loads metadata, but in addition, it selectively stores file instances on disk in accordance with the Disk Retention Rules for written files described in [Selecting Storage Options for a File Group](#).



Files will be restored in the same state as when they were written to the data cartridge. For example, files that were deleted will again appear deleted although visible via the [History Explorer](#).

3.4.4.1 Considerations for LTFS

Tape cartridges formatted with LTFS are interchangeable between XenData6 Server systems with version 6.10 and higher and LTFS systems from third party vendors. The index partition of an LTFS formatted tape is automatically read by the XenData6 Server system and a Volume Catalog is automatically created or updated when a cartridge is imported into a XenData6 Server system.

3.4.4.2 Considerations for ODA

ODA cartridges formatted with the UDF file system are interchangeable between XenData6 Server systems with version 6.20 and higher and with ODA systems from Sony and other third party vendors. The index partition of an ODA cartridge is automatically read by the XenData6 Server system and a Volume Catalog is automatically created or updated when a cartridge is imported into a XenData6 Server system.

3.4.4.3 Considerations for TAR

The **Import Folder Structure** and **Import Data** operations require a Volume Catalog to be present on the cache disk. For a finalized tape, the Volume Catalog is automatically created by the XenData6 Server system when the tape cartridge is imported into the system. Non-finalized tapes do not contain a Volume Catalog and the administrator must use the [Build Catalog](#) operation before using the **Import Folder Structure** or **Import Data** operations.

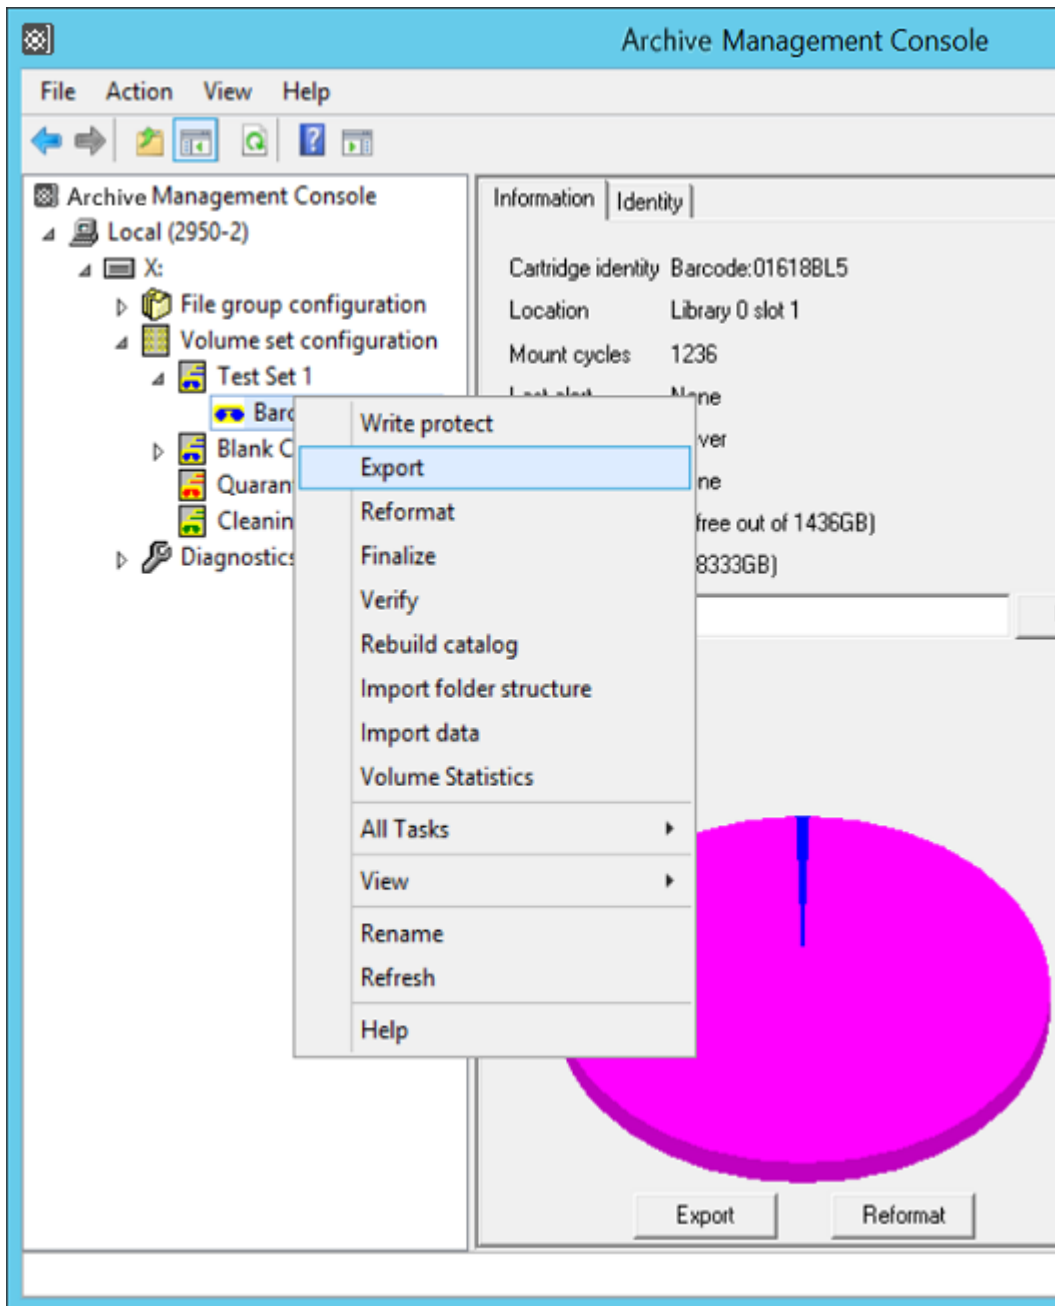
Every time a file or directory is created, renamed or deleted, a record is written both to the tape and to the Volume Contents Catalog. Consider the scenario where a TAR formatted tape is written on one system, transferred to another system where more data is written and then the tape is transferred back to the original system. In that scenario the Volume Contents Catalog on the original system will be out of date and the [Rebuild Catalog](#) operation must be used to update it.

3.4.5 Exporting Cartridges

Data cartridges should be exported from a robotic library using the Archive Management Console as described below and not by using the front panel of the library.

To Export a Data Cartridge

1. Open the Archive Management Console.
2. Navigate to the cartridge.
3. Click **Export** in the right pane of the console or right-click the tape in the left pane and select **Export**.



The selected data cartridge will be moved to one of the robotic library mail slots (often called I/O slots or I/E Elements).

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

3.4.6 Reformatting Cartridges

Reformatting a cartridge erases all the data that is stored on the cartridge and moves the cartridge into the [Blank Cartridge Set](#). When a cartridge is reformatted, files that are stored only on that cartridge (i.e. when there is no replica available and files have been flushed from the cache disk) are made inaccessible. Files that were recorded on the cartridge will be shown in History Explorer as

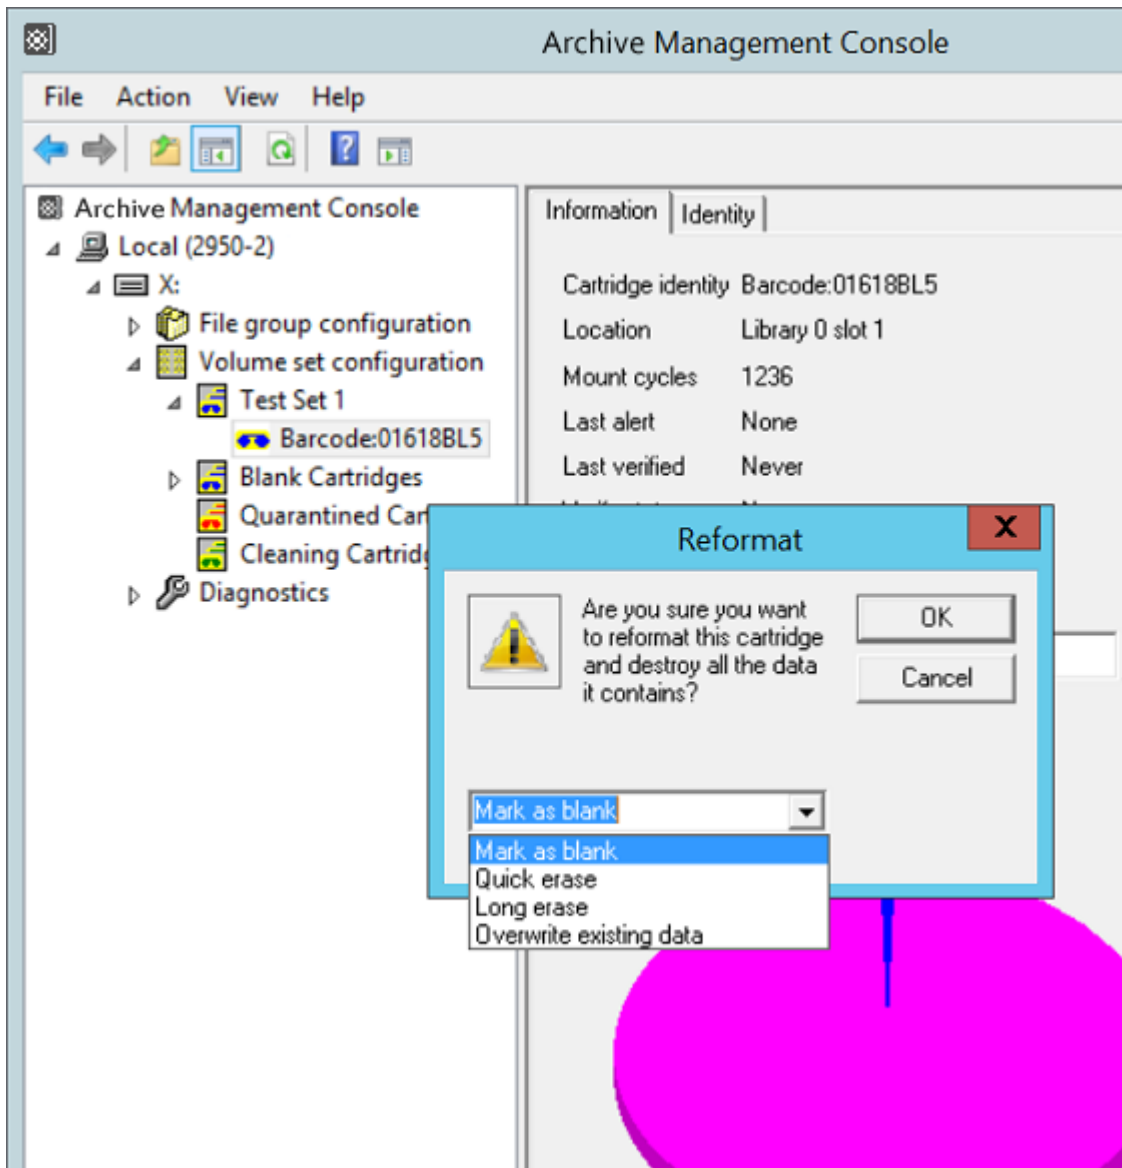
"Archived on unknown volume".

Some 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 Archive Management Console.

To Reformat a Rewritable Cartridge:

1. Open the **Archive Management Console**.
2. Navigate to the cartridge.
3. Click **Reformat**.
4. Choose a reformat option.
 - **Mark as Blank** will instruct the system to treat the cartridge as blank, and move it to the Blank Cartridge Set. However, it will not be overwritten until it is assigned to a new Volume Set. Up to that point, the data may be retrieved by exporting and then re-importing the cartridge.
 - **Quick Erase** writes an end of data mark at or near the beginning of the cartridge (thereby marking the rest of the cartridge blank). This operation does not physically overwrite the data on the cartridge. However, once this operation has been performed it is not possible to recover data 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 existing data** 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 thoroughly overwrites all the data on the cartridge. Used mainly for very sensitive applications.

Note that not all reformat options are supported by all hardware. If a particular option is not visible then your hardware does not support that particular operation.



3.4.7 Displaying Cartridge Information

To Display Information about a Cartridge

1. Open the Archive 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 Volumes that are allocated to the Volume Set.
4. In the case of replicated Volumes, click on the + sign next to the Volume to expand it and show the cartridges that are allocated to the Volume.
5. Select a cartridge to display the information pane, as shown below.

The screenshot displays the Archive Management Console interface. On the left, a tree view shows the navigation structure: Archive Management Console > Local (2950-2) > X: > Volume set configuration > Test Set 1 > Barcode:01618BL5 (selected). The right pane shows the details for the selected cartridge. The 'Information' tab is active, displaying the following data:

Property	Value
Cartridge identity	Barcode:01618BL5
Location	Library 0 slot 1
Mount cycles	1236
Last alert	None
Last verified	Never
Verify status	None
Usage	0.7% used (1426GB free out of 1436GB)
Overhead	0.8% (0.08333GB)

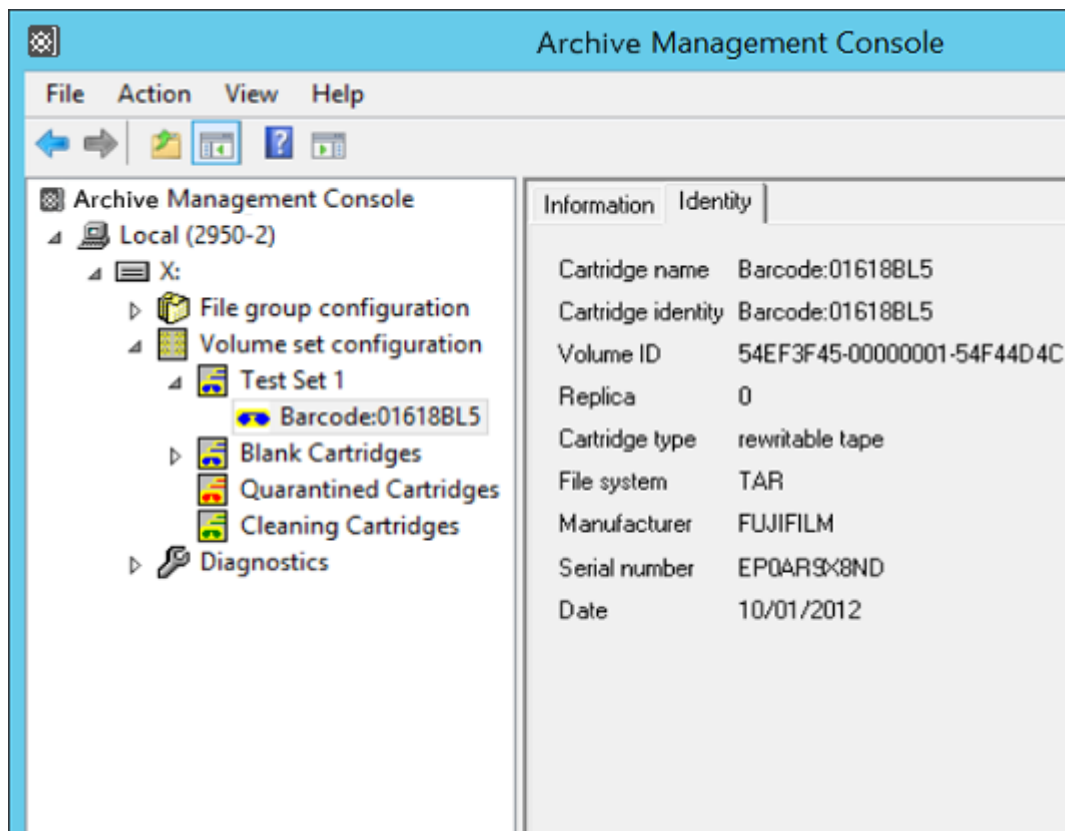
Below the information is an 'Info' field with an 'Update' button. At the bottom of the console is a 3D pie chart showing the space usage: a very small blue segment for used space, a large pink segment for free space, and a small red segment for overhead.

The blue segment in the pie chart represents used space, the pink represents free space and the red represents overhead, which is cartridge capacity that has been consumed but is not used for user data or file system metadata. If the overhead percentage is large, the cause is likely to be one of the following:

- Only a small amount of data has been written to the cartridge - in this case, the overhead percentage is dominated by cartridge format metadata and is not representative of a full cartridge.
- Abnormally high re-writing of data due to the drive error correction - this can be indicative of a faulty drive, a drive that needs cleaning or a faulty cartridge.
- Writing of many mainly small files to the cartridge - there is a fixed overhead associated with creating a file and if there are a large number of small files this can dominate the space consumed on the cartridge.

Note that space consumed by deleted files is not included in the overhead figure.

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

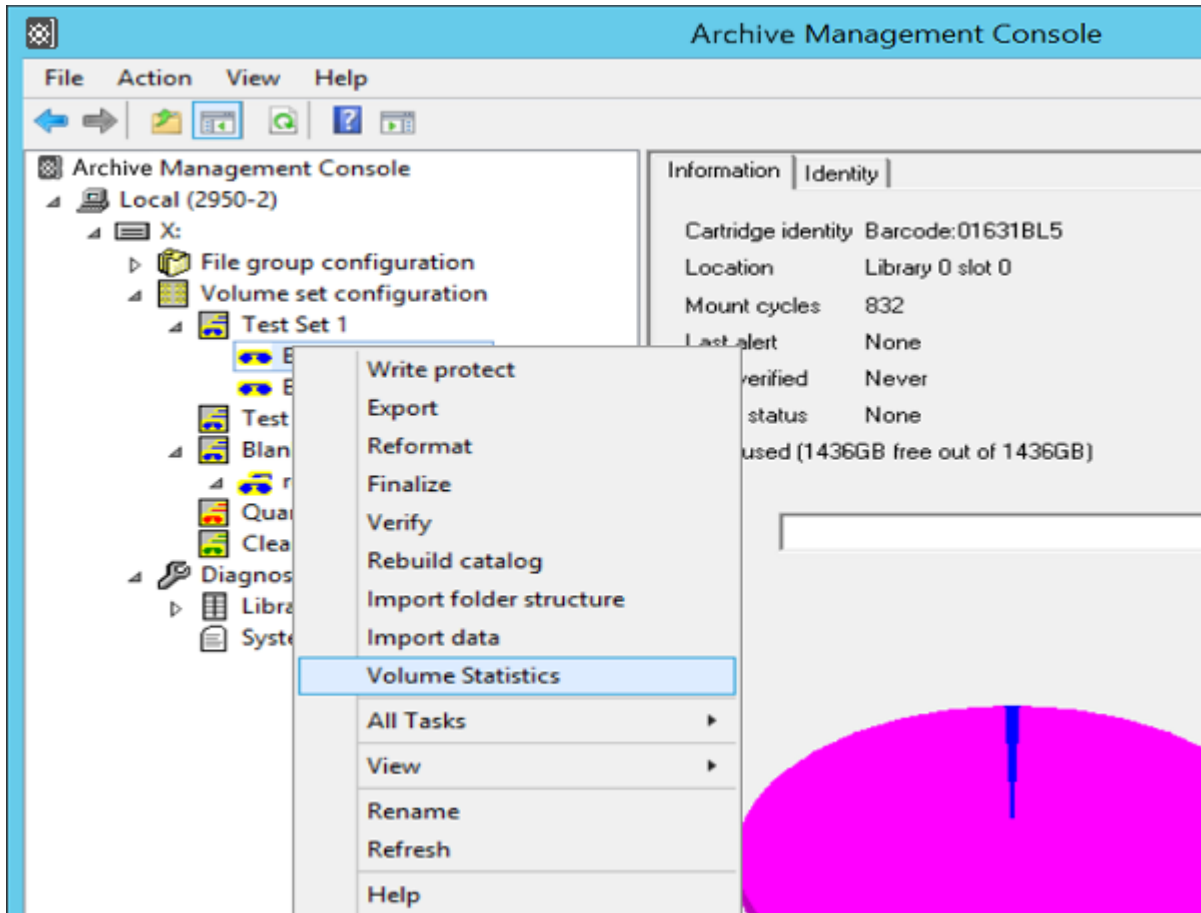


3.4.8 Volume Statistics

The Volume Statistics dialog box gives information about how the space on a Volume is being used. 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 **Archive Management Console**.
2. Navigate to the Volume.
3. Right click on the Volume and select **Volume Statistics**.



An example of Volume Statistics is shown below.

Statistics for Barcode:01618BL5

	Number	Bytes	Bytes on cartridge
Total number of fragment files	50	50.00 GB	50.00 GB
Fragment files currently accessible	50	50.00 GB	50.00 GB
Deleted fragment files	0	0 bytes	0 bytes
Fragment files in old versions of files	0	0 bytes	0 bytes
Rearchived fragment files	0	0 bytes	0 bytes
Fragment files missing metadata	0	0 bytes	0 bytes
Delete and rename records	0	---	0 bytes
Directory records	1	---	16.00 KB
File system overhead			0 bytes
Space that repack would recover			0 bytes

OK Cancel

The statistics provided are:

- Total number of fragment files - the total number of fragment files and the amount of space they occupy on the Volume. When fragmentation is not enabled, this is the same as the total number of files.
- Fragment files currently accessible - information about files that are accessible via the Windows file system (excluding old versions of files and deleted files).
- Deleted fragment files - information about deleted files and the space they occupy.
- Fragment files in old versions of files - information about old (overwritten) files and the space they occupy.
- Re-archived fragment files - for a Volume that has been partially repacked, this gives information about files that have been transferred to other Volumes using a Repack operation.
- Fragment files missing metadata - files that are stored on the Volume but do not have complete file system metadata on the cache disk. The **Import Folder Structure** operation can be used to update the metadata.
- Delete and rename records - these are records that are written to a Volume when a file is deleted or renamed.
- Directory records - these are records that are written to a Volume when a folder is created or deleted.
- File system overhead - this is space that is consumed by the file system that is not accounted for by the other entries in this list.
- 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 Archive 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 also be exported for use in Excel and other programs.

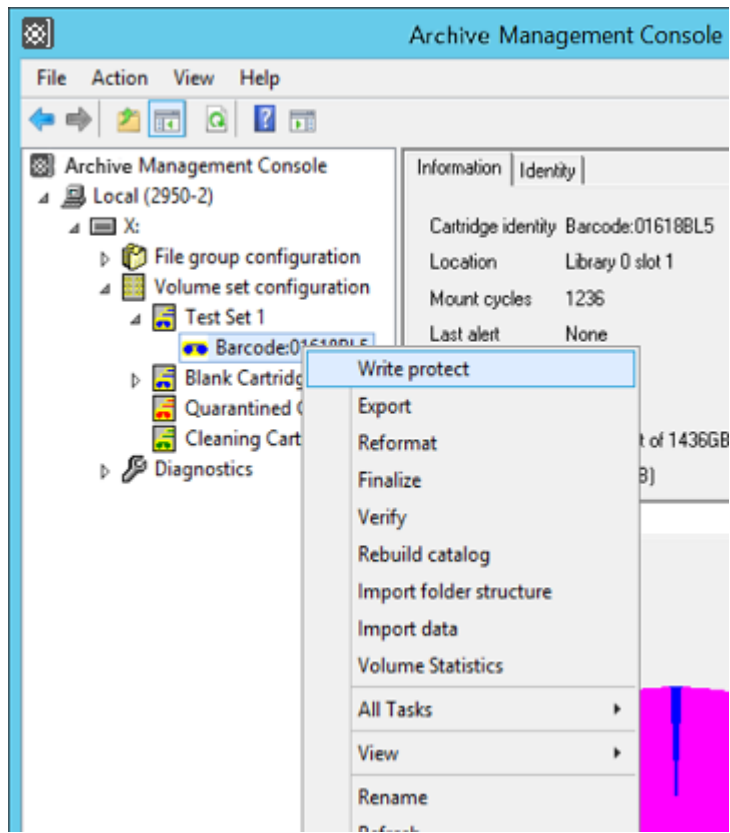
3.4.9 Write-Protecting Cartridges

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

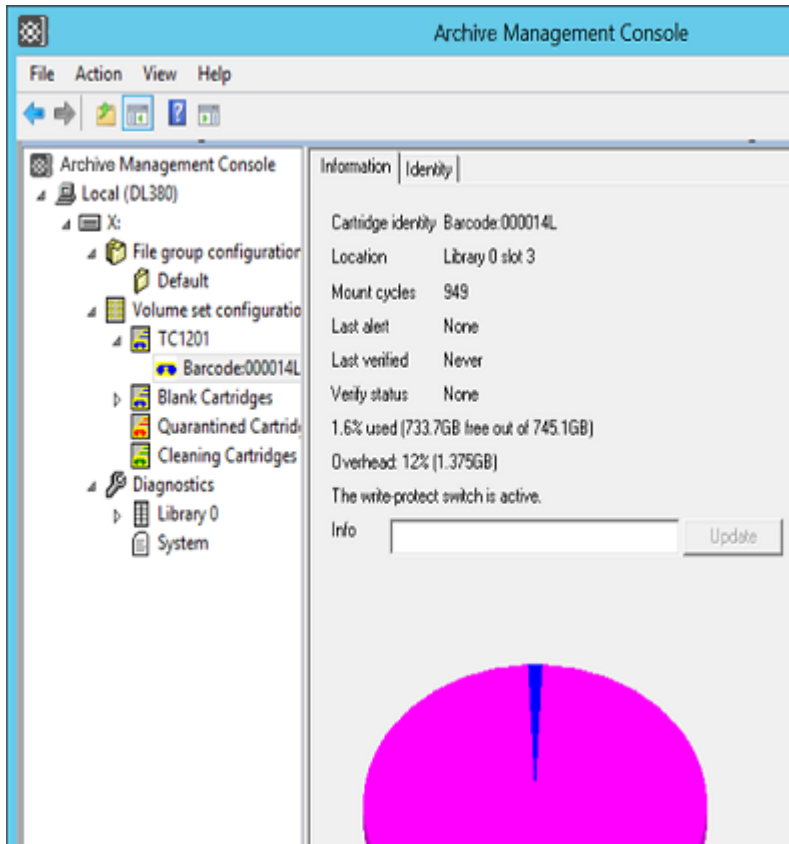
This can be achieved by write protecting the cartridge (or replicated set of cartridges that form a Volume). If all the Volumes in a Volume Set are full, finalized or write-protected, the administrator will have to add a new Volume before more data can be written to the Volume Set.

To Write-Protect a Data Cartridge

1. Open the Archive Management Console.
2. Navigate to the cartridge.
3. Right click on the cartridge and select **Write protect**.



In addition to the software write-protection described above, a cartridge can be write protected by using the physical write protect switch on the cartridge. When the cartridge write protection switch is active, this will be identified on the [information page](#) of the medium properties for that cartridge, as shown below.



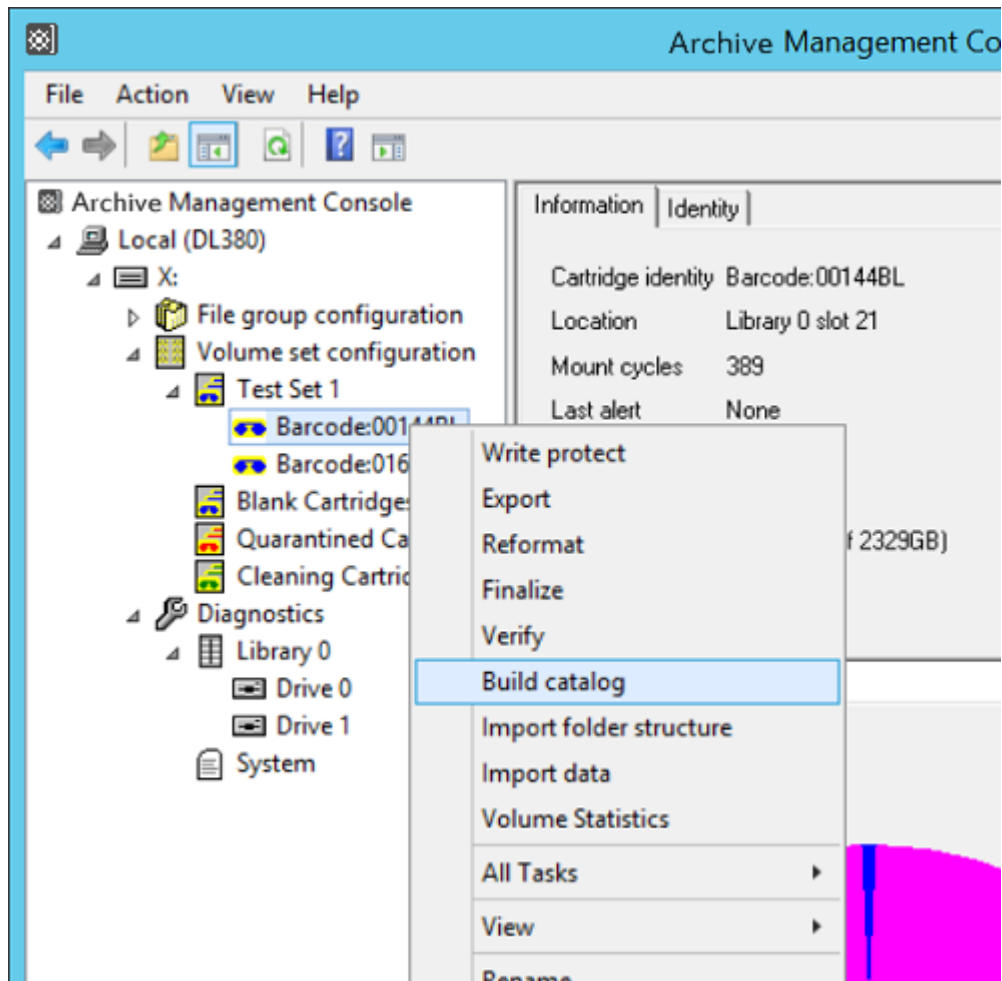
3.4.10 Building and Rebuilding Volume Contents Catalogs

A [Volume Contents Catalog](#) is a file on the cache disk that contains a journal of the history of a Volume and is useful in many situations. Under some circumstances, a cartridge that is known to the system may not have a Volume Contents Catalog, or the Catalog may be out of date or unusable. For example, this will occur when a non-finalized TAR formatted tape is transferred from another XenData6 Server system.

When a Volume Catalog does not exist on the cache disk, the system provides a mechanism to build it as follows:

To Build the Volume Catalog for a Cartridge

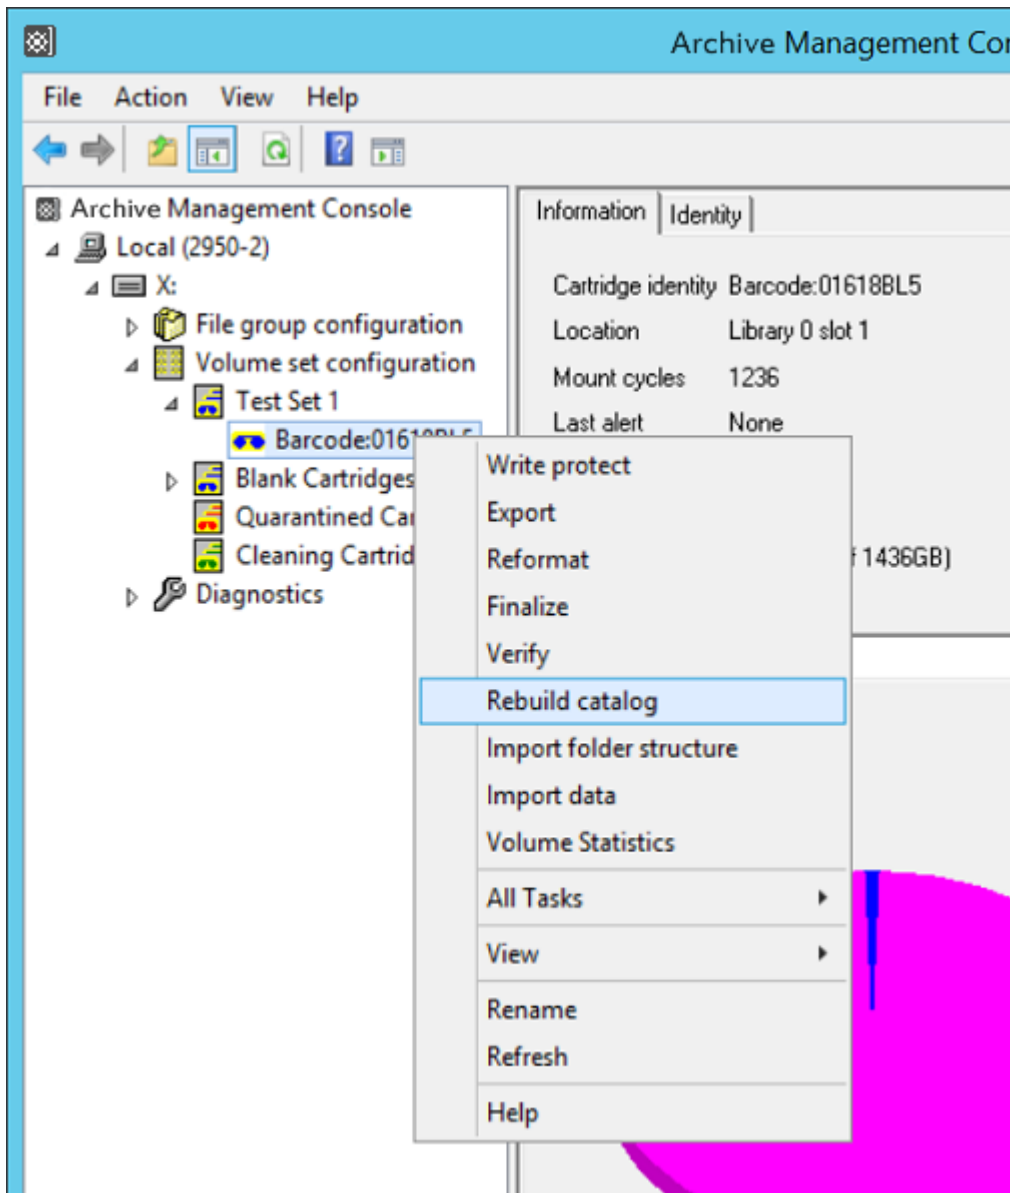
1. Open the Archive Management Console.
2. Navigate to the Volume or cartridge.
3. Right click on the Volume or cartridge and select **Build Catalog**.



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

To Rebuild the Volume Catalog for a Cartridge

1. Open the Archive Management Console.
2. Navigate to the Volume or cartridge.
3. Right click on the Volume or cartridge and select **Rebuild Catalog**.



3.4.11 Repacking Volumes

Repack is an operation which copies files from one Volume to another, omitting 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. It can be used to recover space that is consumed by old versions of files that have been overwritten or deleted, and it is also used to move data from one cartridge format to another (for example when a new, higher capacity cartridge format becomes available).

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 target destinations defined by the current [File Group rules](#). A File Group rule must exist for all files that are stored on the Volume that is being repacked.
- When all the files on the Volume have been successfully repacked, the repacked cartridges are moved to the [Quarantined Cartridge Set](#) 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.

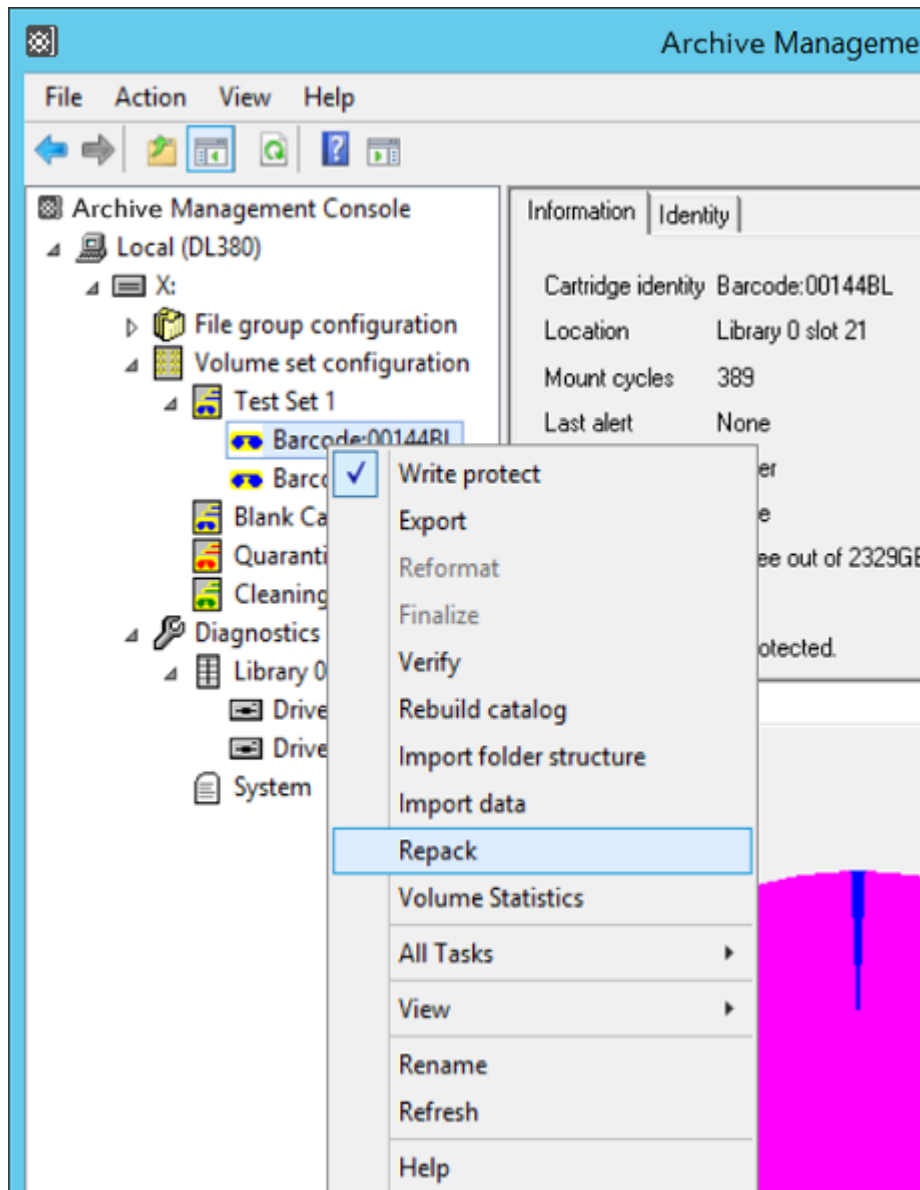
Hardware Requirements

We recommend that the repack operation should be performed on an archive system with at least two drives. The repack operation cannot be performed on an archive with no library and only one stand-alone drive, but it can be performed on an archive system with a robotic library having only one drive.

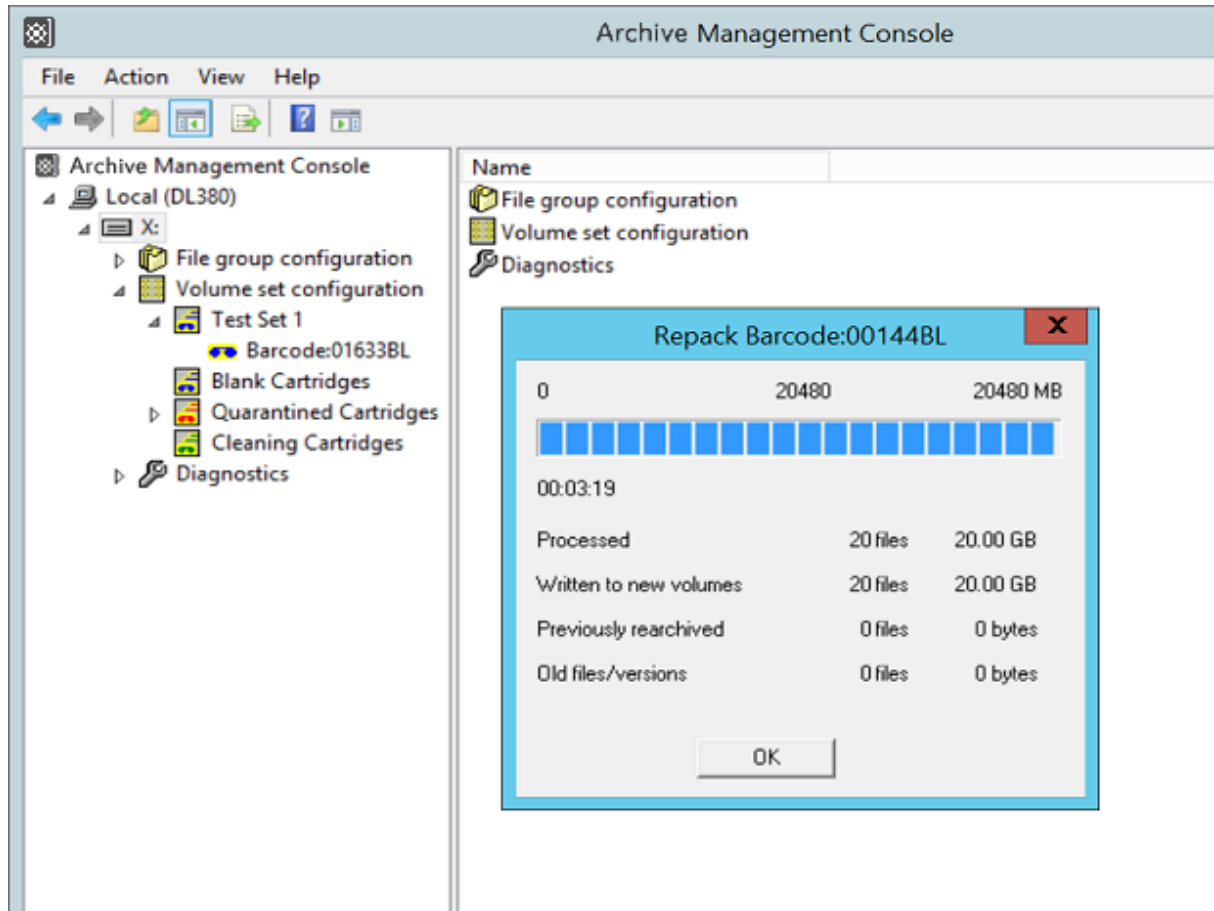
Note that this might be a slow operation if the amount of free space on the cache disk is less than the amount of data being repacked.

To Repack a Volume

1. Open the Archive Management Console.
2. Navigate to the Volume.
3. Ensure that the Volume is not writable. If it is writable, [write-protect](#) it.
4. Right click on the Volume and select **Repack**.



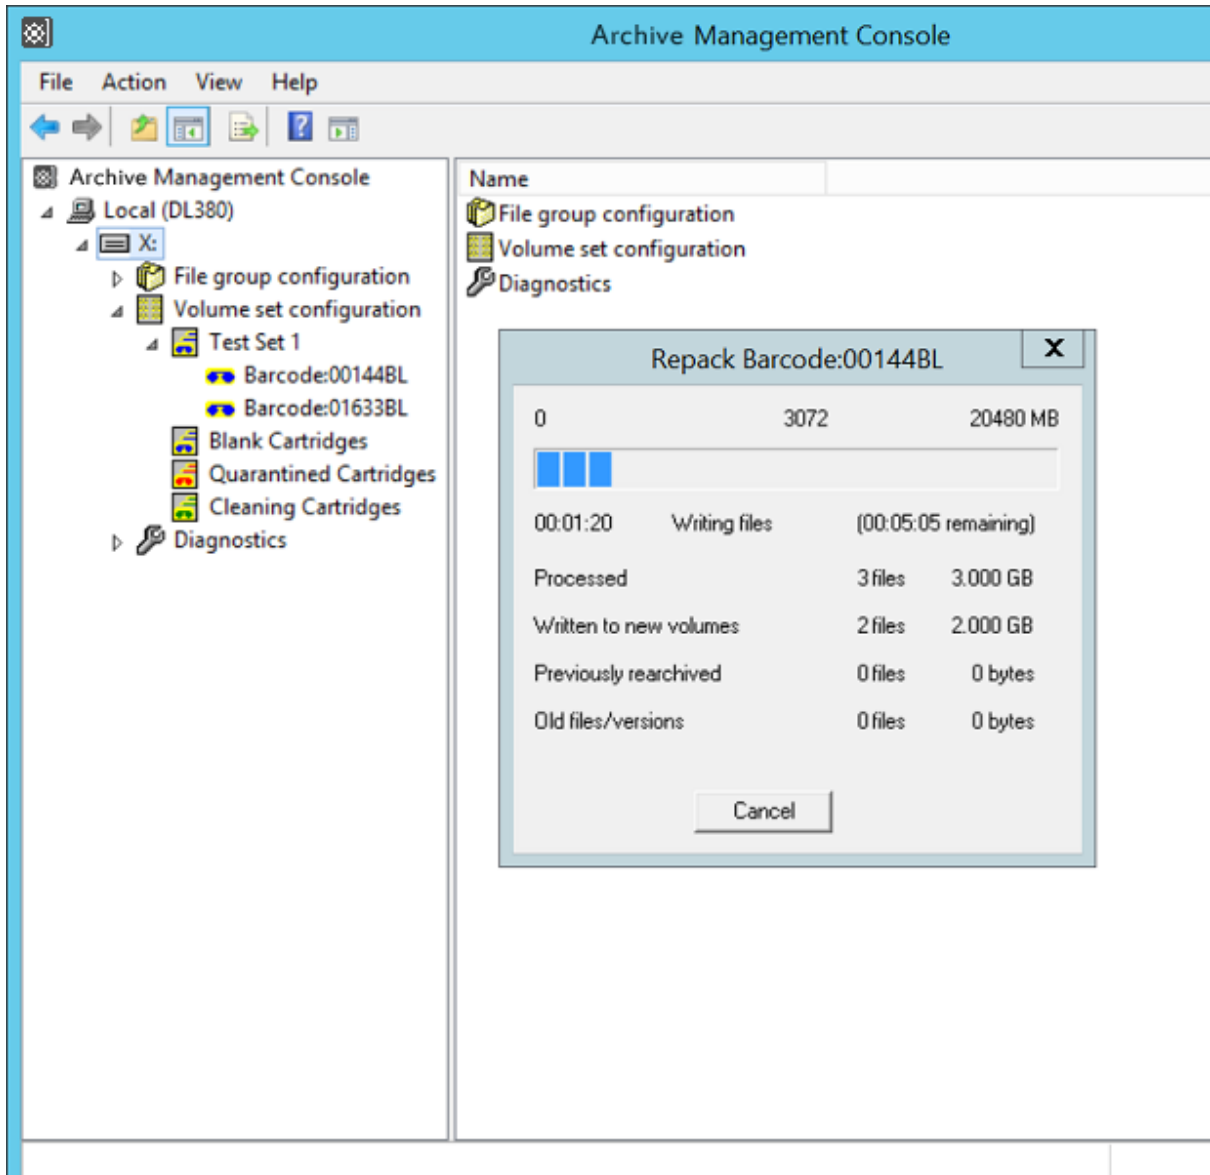
When the repack operation has been successfully completed, the cartridges from the repacked Volume will be moved to the Quarantined Cartridge Set and status information as shown below will be displayed. Quarantined rewritable cartridges may then be [reformatted](#) for re-use.



3.4.12 Canceling a Repack Operation

The Repack operation, described [here](#) may take many hours to complete. The operation can be canceled and restarted at a later time. When the repack operation is running, a progress box is displayed as shown below and the operation can be canceled by clicking **Cancel**.

Note that if a canceled Repack operation is restarted at a later time, it will resume from where it was previously canceled.

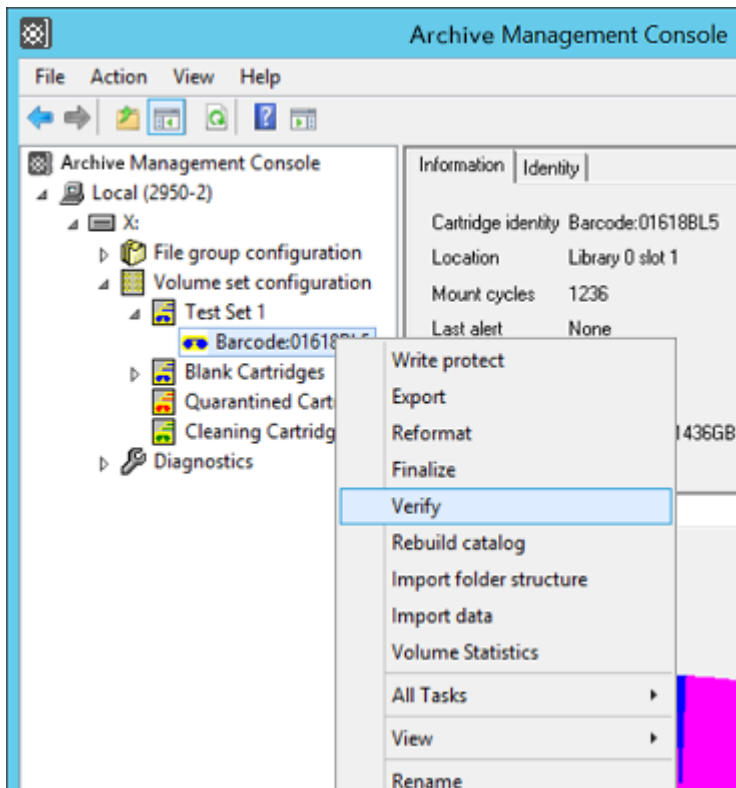


3.4.13 Verifying the Data on a Cartridge

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

To Verify the Data on a Cartridge

1. Open the Archive Management Console.
2. Navigate to the cartridge.
3. Right click on the cartridge and select **Verify**.



3.4.14 Finalizing Volumes

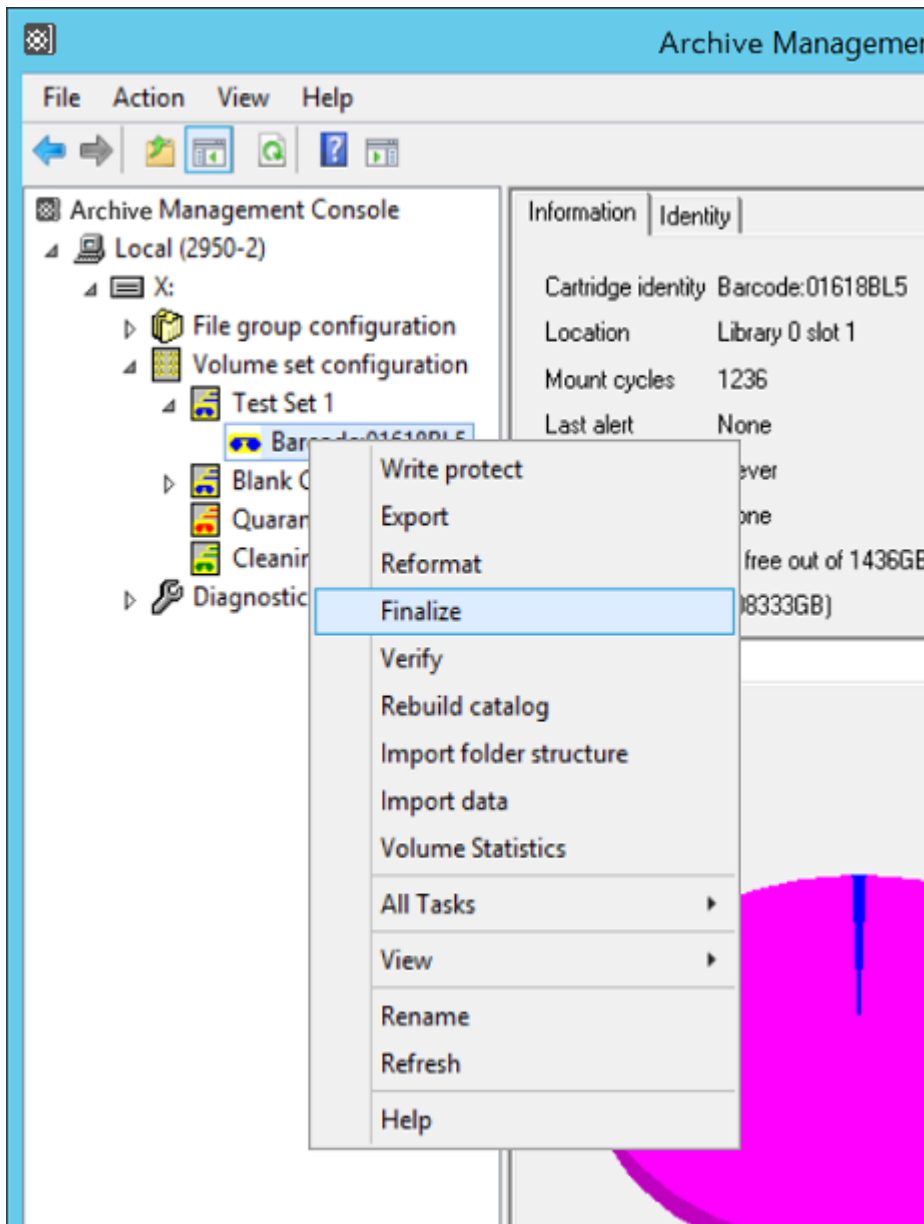
The [Finalization](#) operation is available for tape cartridges formatted using the TAR file system and for ODA WORM cartridges. Finalizing a cartridge results in special marks being written that prevent additional data being written to the cartridge. In addition, for TAR format tapes that have a Volume Contents Catalog, the Finalization operation writes the Catalog to the end of the tape.

When an eligible Volume becomes 98% full, the system automatically finalizes it. Volumes may be finalized before they are 98% full as described below.

To Finalize a Volume

Open the Archive Management Console.

1. Navigate to the Volume.
2. Right click on the Volume and select **Finalize**.

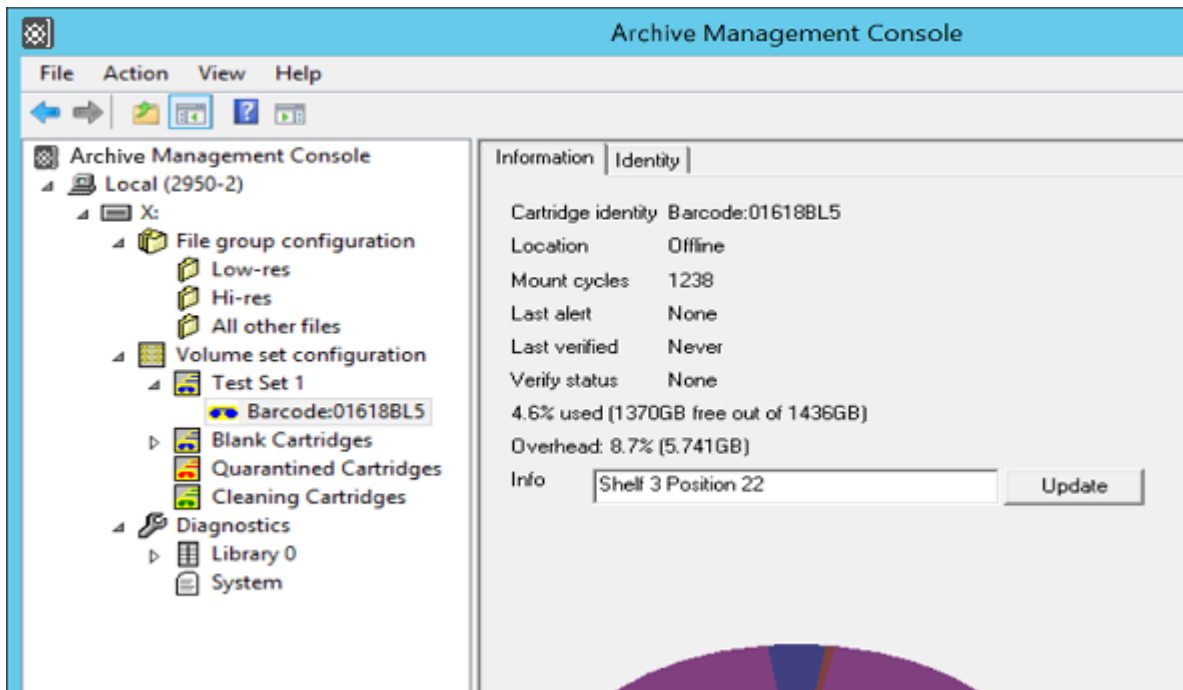


3.4.15 Setting the Administrator Defined Information for a Cartridge

The Archive Management Console contains an administrator-defined information field for every cartridge. This field can contain any text up to 63 characters in length and is typically used to describe the offline location of a cartridge (i.e. "Shelf 1") or to describe the contents of the cartridge.

To Set the Administrator Defined Information for a Cartridge

1. Open the **Archive Management Console**.
2. Navigate to the cartridge.
 1. Enter the desired information in the "Info" box.
3. Click **Update** to save your changes.



Using Mac Clients

4 Using Mac Clients

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

When a Mac client computer running OS/X writes resource forks and extended attributes to a Windows SMB share (such as a XenData6 Server archive) it writes either AppleDouble files or files with [Alternate Data Streams](#). This version of XenData6 Server software fully supports AppleDouble files and manages them using hidden File Group rules. It also fully supports alternate data streams when writing to tape cartridges using the TAR format and provides limited support for common alternate data streams when writing data in the LTFS or ODA formats.

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 are fully supported by the TAR file system and partially supported by other file systems, and are managed by hidden File Group rules.

In addition, there are two areas of file management which require configuration of each Mac client that has access to the XenData6 Server system:

- Icon Preview on any application running on the Mac must be disabled for all Mac clients that access the archive, otherwise unwanted file restores will result which can greatly impair the archive's performance; and
- Alternate Data Streams will need to be disabled for Mac clients that write and restore files to tape cartridges formatted with LTFS or to ODA cartridges.

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

4.1 Hidden File Group Policies

In addition to the File Group rules defined by the administrator using the Archive Management Console, hidden File Group rules are automatically implemented for improved management of certain types of file including 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 files to store custom attributes of a folder such as background color and position of icons. The XenData6 Server system will store any file named .DS_Store on the disk cache but will not save the file to the archive. This rule overrides all policies defined in the Archive Management Console.

When a Mac client computer running OS/X writes resource forks and extended attributes to a Windows SMB share (such as a XenData6 Server archive) it writes either AppleDouble files or files with [Alternate Data Streams](#). 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 be named 'abc.mov' and the hidden resource fork file will be named '._abc.mov'. With a [Hierarchical Storage Management](#) system, it is important that resource fork files are not flushed from the disk cache as they are accessed frequently by Mac clients. The hidden file group policies prevent flushing of resource fork files. In cases where the corresponding data file is saved to a Volume Set, the resource fork will also be saved to the same volume set. These rules override all policies defined in the Archive Management Console.

When OS/X writes a file with an alternate data stream, the alternate data stream is usually very small; it contains the same information as is held in the resource fork part of an AppleDouble file. As for resource forks, the alternate data stream is accessed frequently by Mac clients and consequently it is important that alternate data streams are not flushed. The XenData6 Server system does not flush alternate data streams.

The TAR tape file system is the only file system used by XenData6 Server that supports writing of alternate data streams to data cartridge Volumes. The other file systems (LTFS and ODA/UDF) do not support alternate data streams and will raise an error if an attempt is made to write files that contain them. In order to maximize the general flexibility of the system, it implements default file group rules for some common alternate data streams that are considered non-essential. The default rules store the alternate data streams on the cache disk but do not attempt to write them to data cartridge Volumes, regardless of the format used. The alternate data streams that are treated this way are as follows:

```
:Zone.Identifier
:AFP_AfpInfo
:com.apple.metadata*
:com.apple.quarantine
:com.apple.TextEncoding
:com.apple.FinderInfo
```

Note: Attempts to write any other alternate data streams to the archive will result in the alternate data stream being written to the Volume Set specified for the main file in the Archive Management Console. If the Volume Set specifies writing to a cartridge that is not formatted with the TAR tape file system this will result in an error.

4.2 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.

4.3 Disabling Alternate Data Streams

From OS/X version 10.6, Apple uses alternate data streams as the default configuration rather than writing to a XenData6 Server SMB share (previously, AppleDouble files were the default). However, the XenData6 Server system has limited support for alternate data streams when writing to LTFS formatted tape cartridges and ODA cartridges. In these cases it may be desirable to disable the use of alternate data streams by clients writing to a XenData6 Server archive.

4.4 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 system supports OS/X usage of folders and files that contain these Mac specific characters. Files and folders containing Mac specific characters written to the XenData6 Server archive share are seen by Mac users as they were created.

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

Windows Explorer Extensions

5 Windows Explorer Extensions

XenData6 Server software extends the capabilities of Windows Explorer to provide the following functionality:

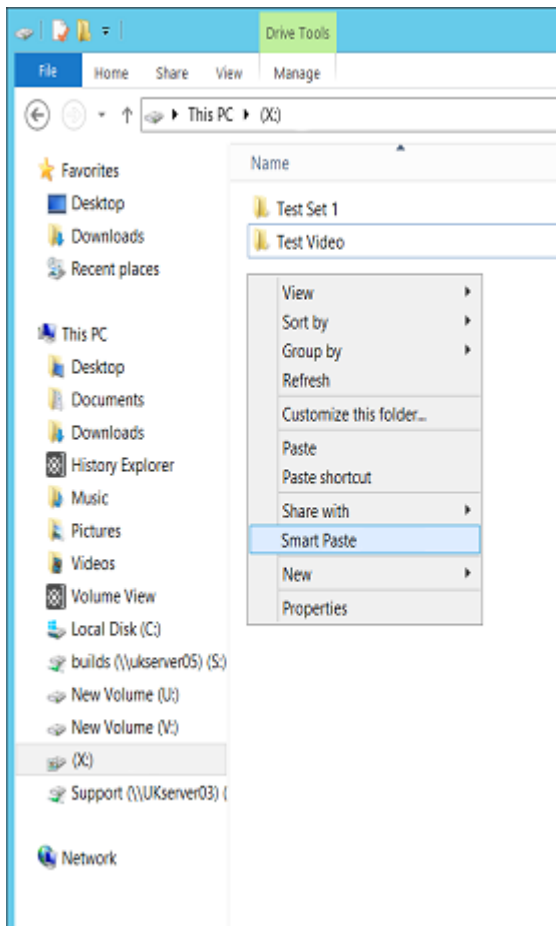
- [Smart copy and paste for optimized file restores](#)
- [Flushing of files and folders](#)
- [Pre-fetching of files and folders](#)
- [History Explorer](#)
- [Volume View](#)
- [Enhanced properties](#)

5.1 Smart Copy and Paste

The standard copy and paste operations available within Windows Explorer restore files in an order which does not take into account the location of the files on data cartridges. When multiple files are being restored, this can cause considerable delays due to excessive cartridge swap operations and non-optimal restore order of files within an individual cartridge. The 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 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 **Smart Paste**.



To Restore Files using Smart Copy

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

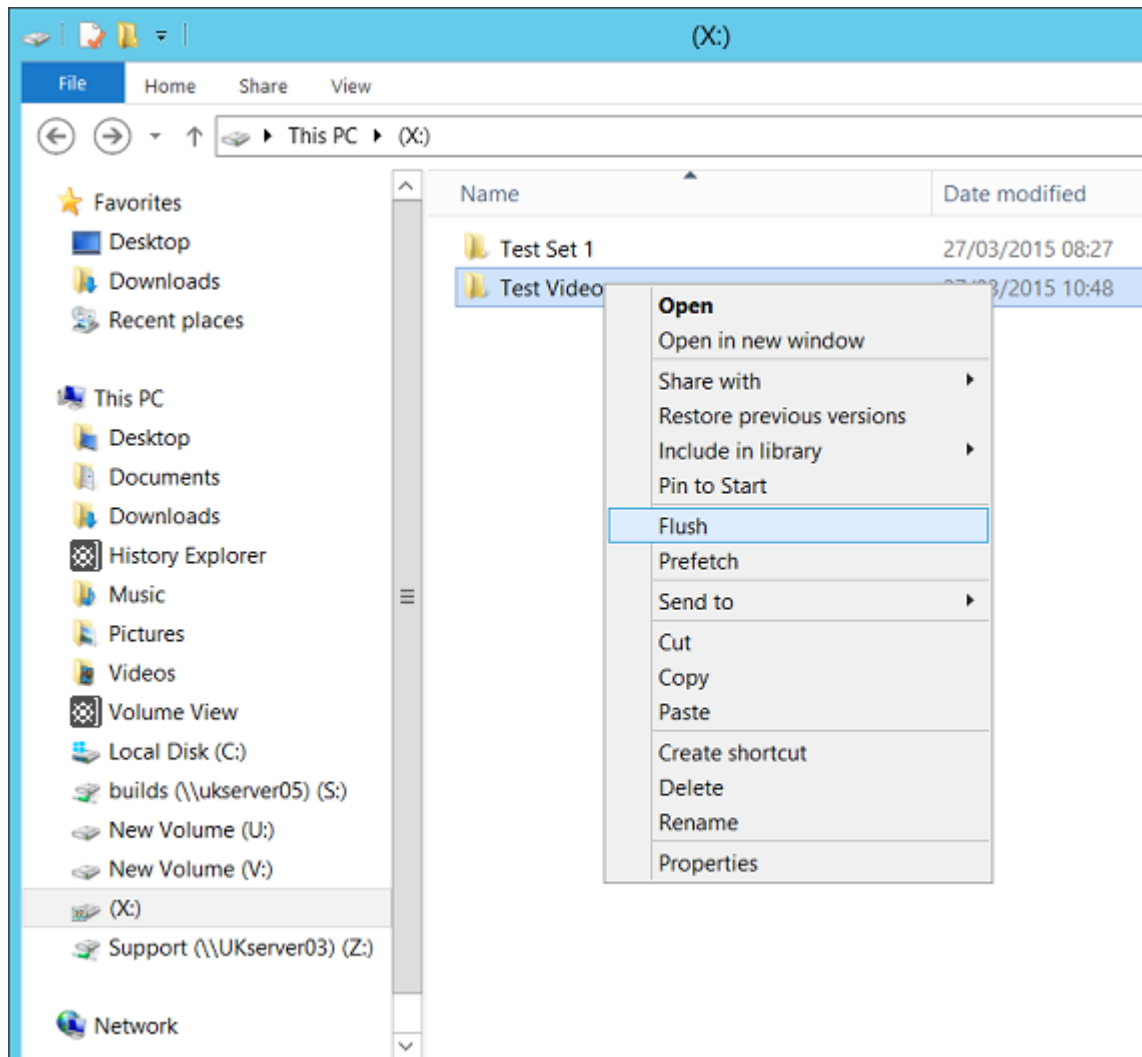
5.2 Flushing

Selected files and the contents of selected folders can be flushed from the cache disk using the Windows Explorer **Flush** option. Flushing will only occur for files that have been successfully written to the designated number of replica cartridges. The Explorer Flush option overrides the [Disk Retention Rules](#) defined in the Archive 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 data to be removed from the cache disk, but the file is still visible and accessible to normal applications by restoring from near line data cartridge Volumes. The Windows offline attribute is set for all files that have been flushed.

To Flush Files using Windows Explorer

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



Note: Windows Explorer sometimes spontaneously reads files after a flush operation. If the applicable disk retention rules defined in the Archive Management Console are not set to flush immediately after a file is closed, this will result in the file being fetched back to the cache disk.

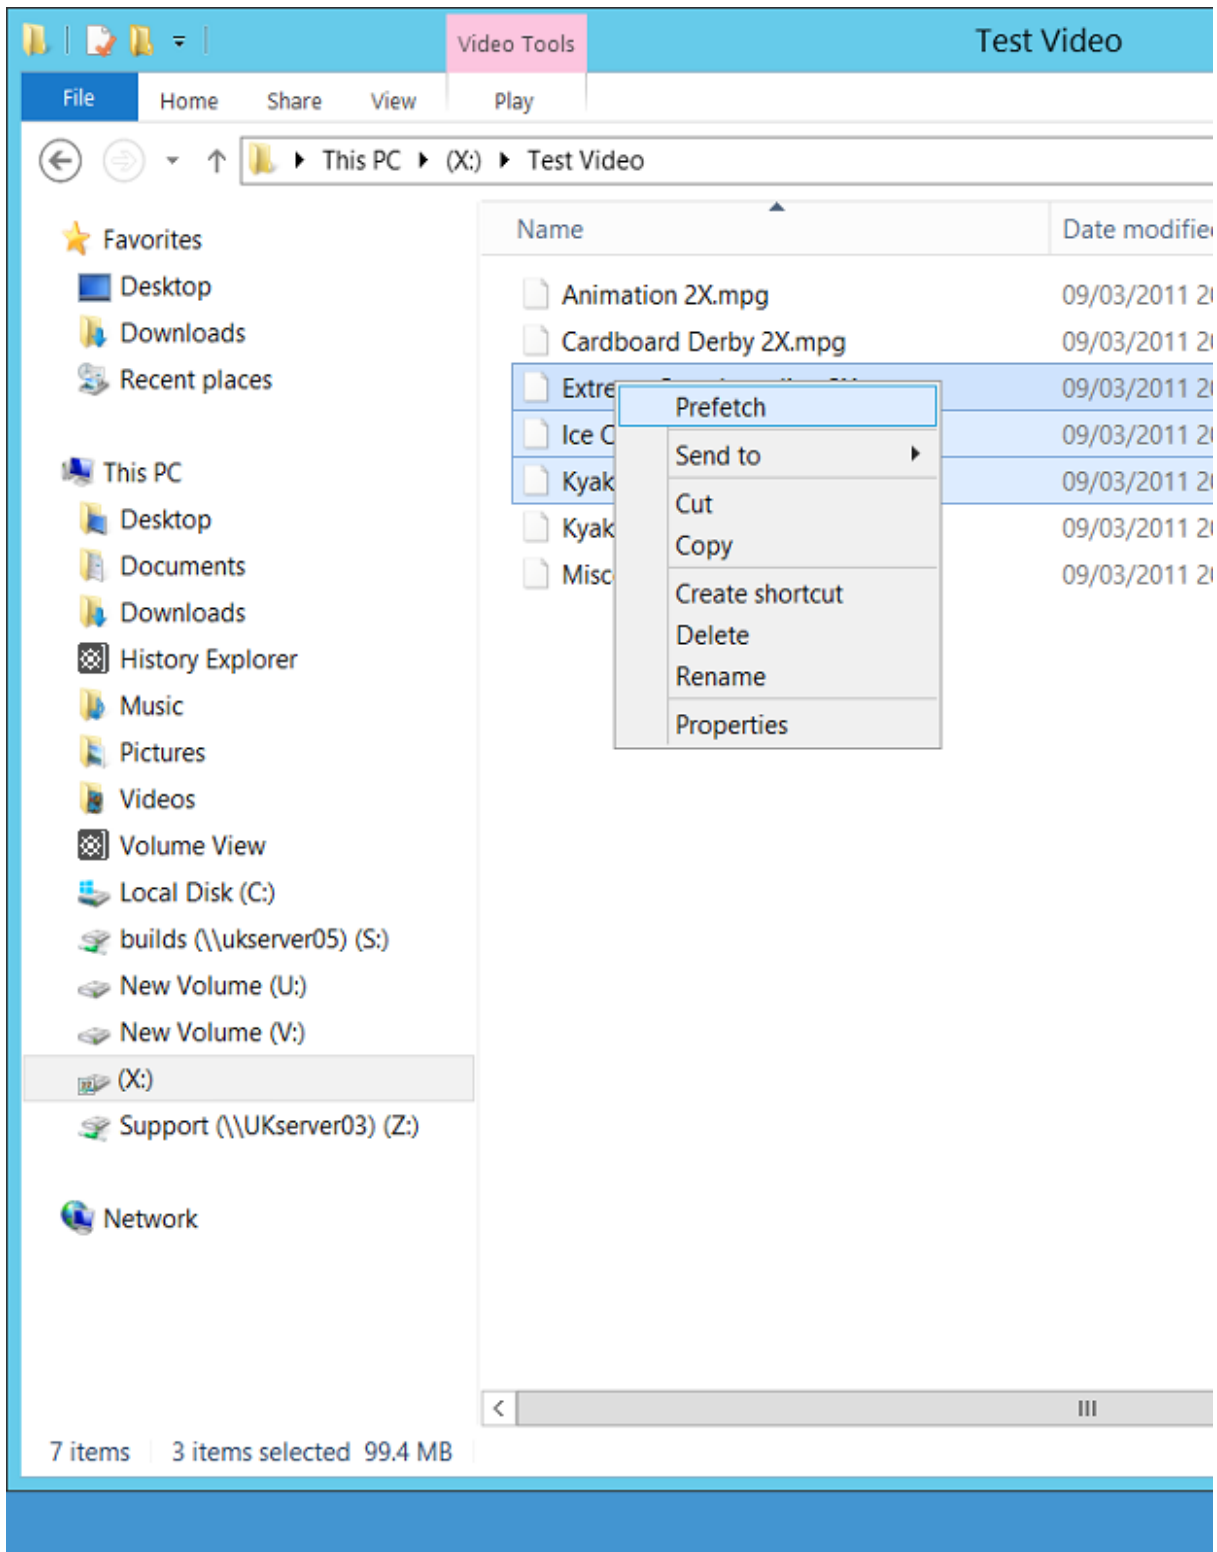
5.3 Pre-fetching

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

The Explorer Prefetch option overrides the [Disk Retention Rules](#) defined in the Archive Management Console. Pre-fetched files will remain on the cache disk until they have been read (when the **Flush read files from disk** Retention Rule will be applied) or until they are manually [Flushed](#) using Windows Explorer.

To Prefetch Files using Windows Explorer

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



Windows Explorer sometimes spontaneously reads files after a pre-fetch operation. If the disk retention rules defined in the Archive Management Console are set to flush after a file is closed, this will result in this file being flushed from the cache disk.

Note that if only a single file is selected, a **Recall** option is also available. This is very similar to the

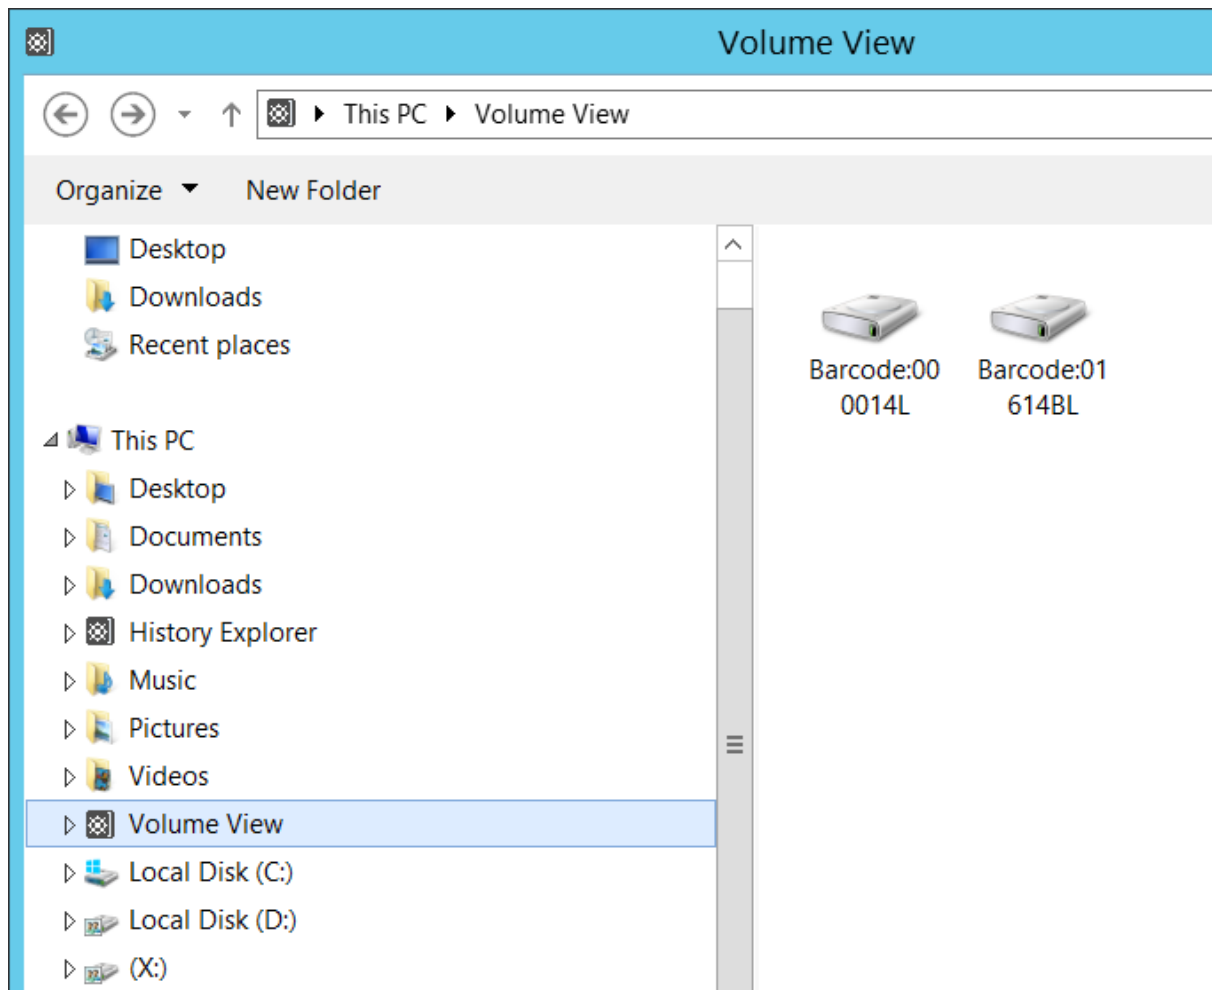
Prefetch operation but additionally provides an on-screen display of any applicable error messages.

5.4 Volume View

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

To Browse with Volume View using Windows Explorer

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



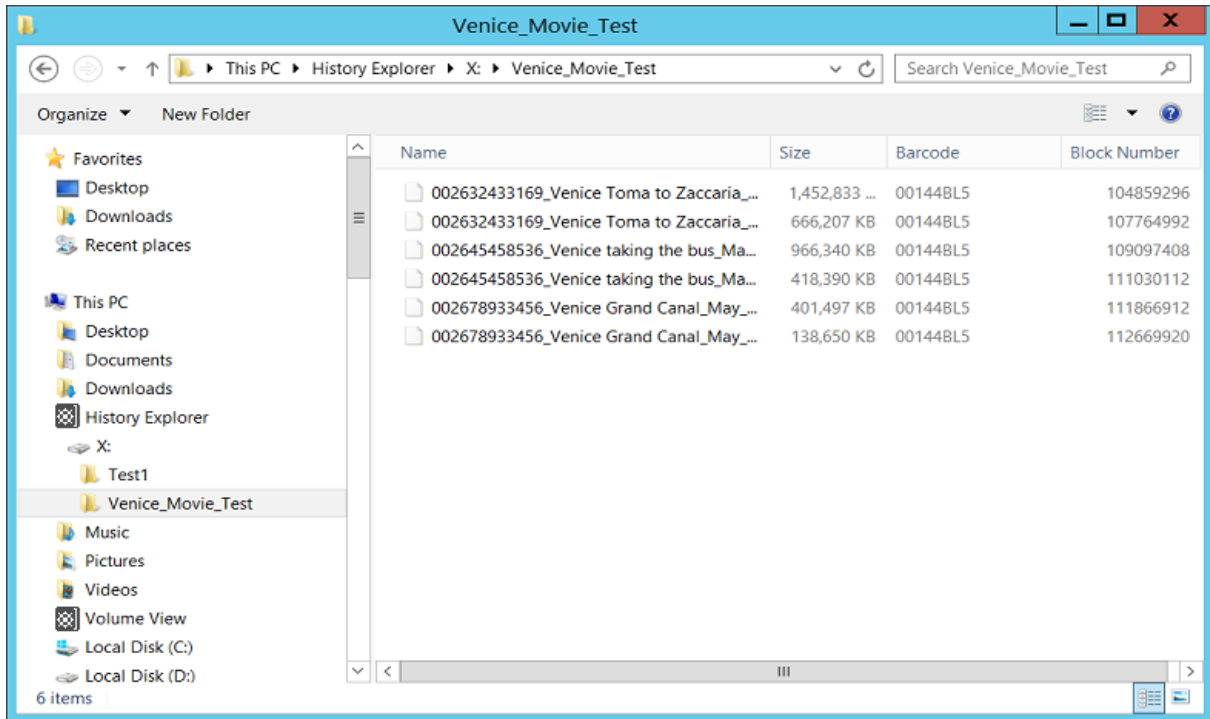
5.5 History Explorer

The History Explorer is used to obtain the complete history and status of any file that the system knows about. The History Explorer lists all available versions of all files, all file instances and their cartridge locations, including deleted and renamed files. It also allows the retrieval of old, overwritten or deleted file versions.

To Browse with History Explorer

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

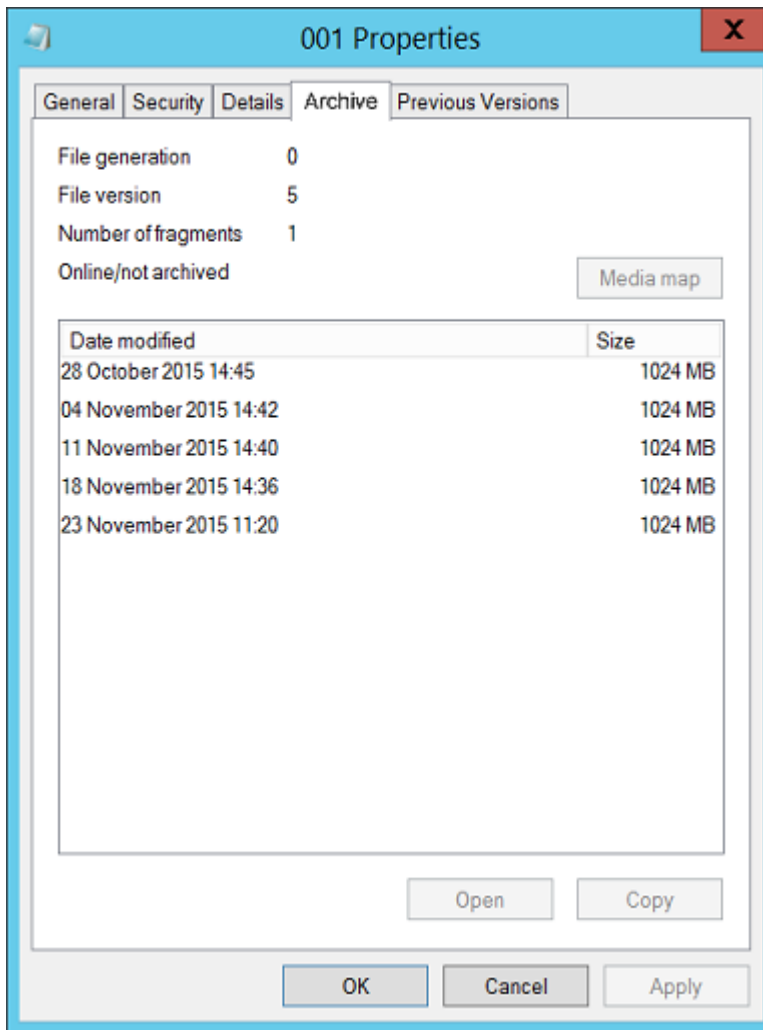
Deleted files are shown with a grayed out icon:



XenData6 Server software adds an additional tab to a file's properties dialog. The Archive 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. Open Windows Explorer.
2. Select and then right-click on the required file.
3. Select **Properties**.
4. Select the **Archive** tab.



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**.

Note: the **Open** option is not available for all file types because not all applications support the required interaction with the XenData6 Server Windows Explorer extension.

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

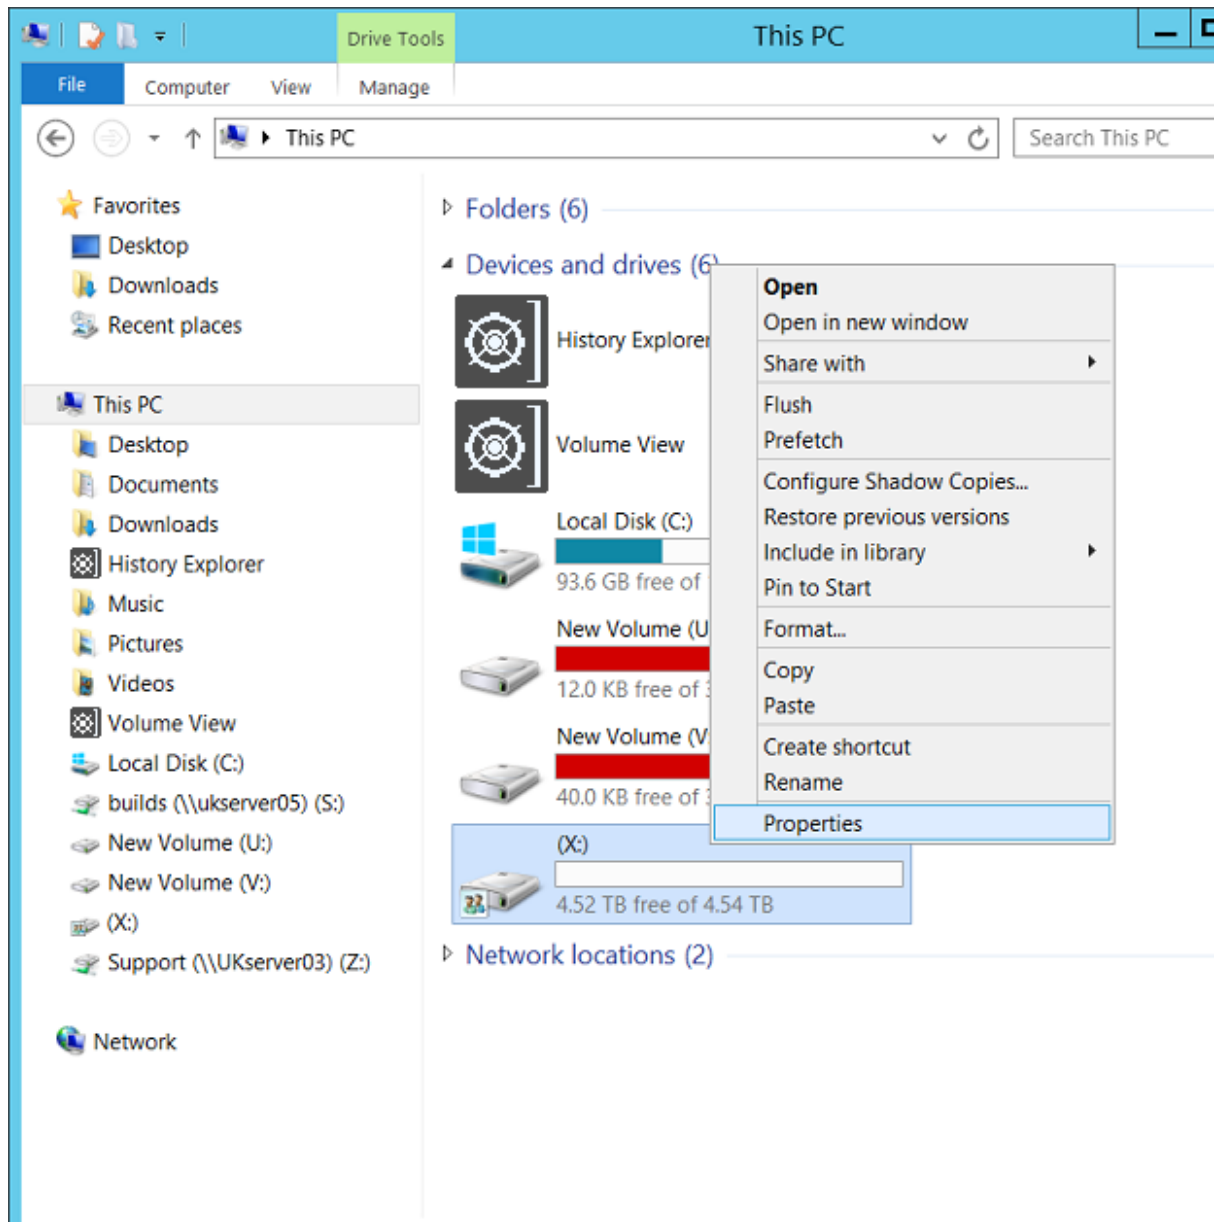
1. View the file versions as described above.
2. Click on the required version of the file.
3. Click **Copy**.
4. Use Windows Explorer to paste the file to the required storage location.

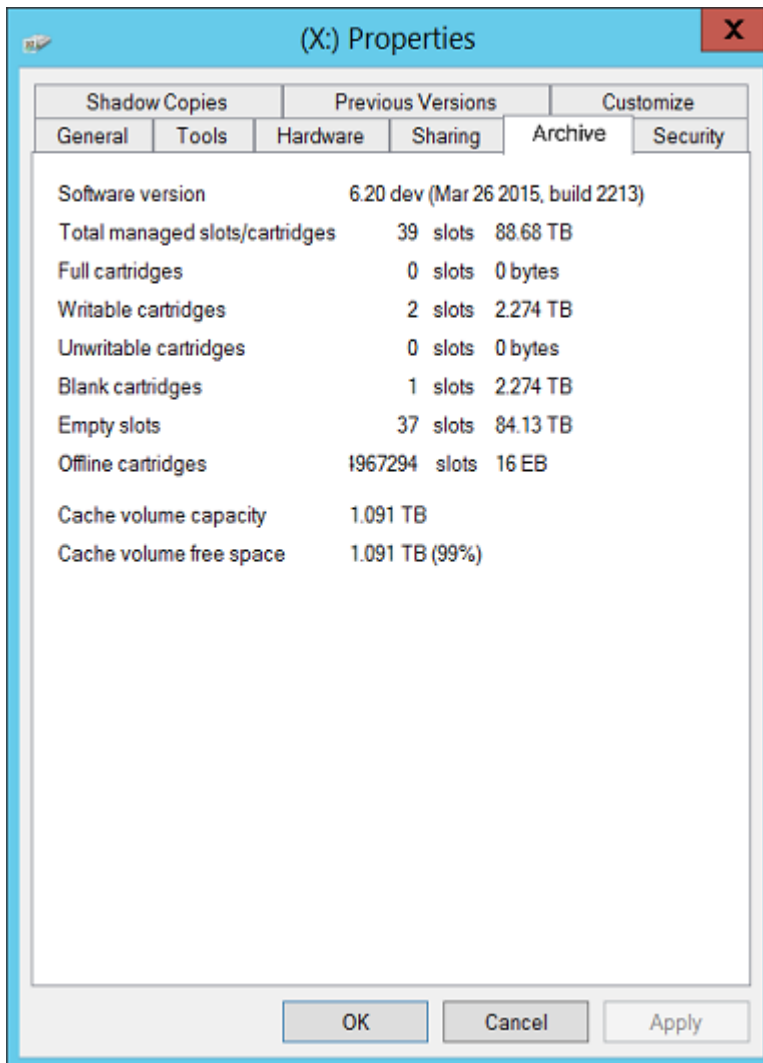
5.6 Enhanced Properties for the Archive Volume

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

To Obtain Enhanced Properties

1. Open Windows Explorer
2. Right click on the logical drive letter under control of XenData6 Server.
3. Select **Properties** and then select the **Archive** 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 robotic libraries and stand-alone drives plus the number of managed offline cartridges.
- Full cartridges - the total number of full and finalized cartridges including offline cartridges that are full or finalized and cartridges that are only partially filled but have been manually finalized.
- 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 robotic libraries and drives.
- Empty slots - the number of empty slots in attached robotic libraries and the number of empty stand-alone drives.
- Offline cartridges - the number of offline cartridges managed by the system.
- Cache volume capacity - the capacity of the cache disk volume
- Cache volume free space - the free space available on the cache disk volume.

Metadata Backup

6 Metadata Backup

The Metadata Backup program deals with:

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

If a XenData6 Server system has to be rebuilt, perhaps due to a RAID failure, the file system metadata may be rebuilt by using the **Build Catalog**, **Import Folder Structure** and **Import Data** functions available in the Archive Management Console. However, this can be a lengthy process for a system with a large number of Volumes. The Metadata Backup program speeds up the process of rebuilding the data on the cache disk by restoring the file system metadata and State File to the condition they were in at the time of the metadata backup. This means that the Build Catalog and Import Folder Structure functions need only be used for Volumes which have been written since the latest backup.

6.1 Starting Metadata Backup

For Windows Server 2008 R2 and Windows 7:

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

For Windows Server 2012R2:

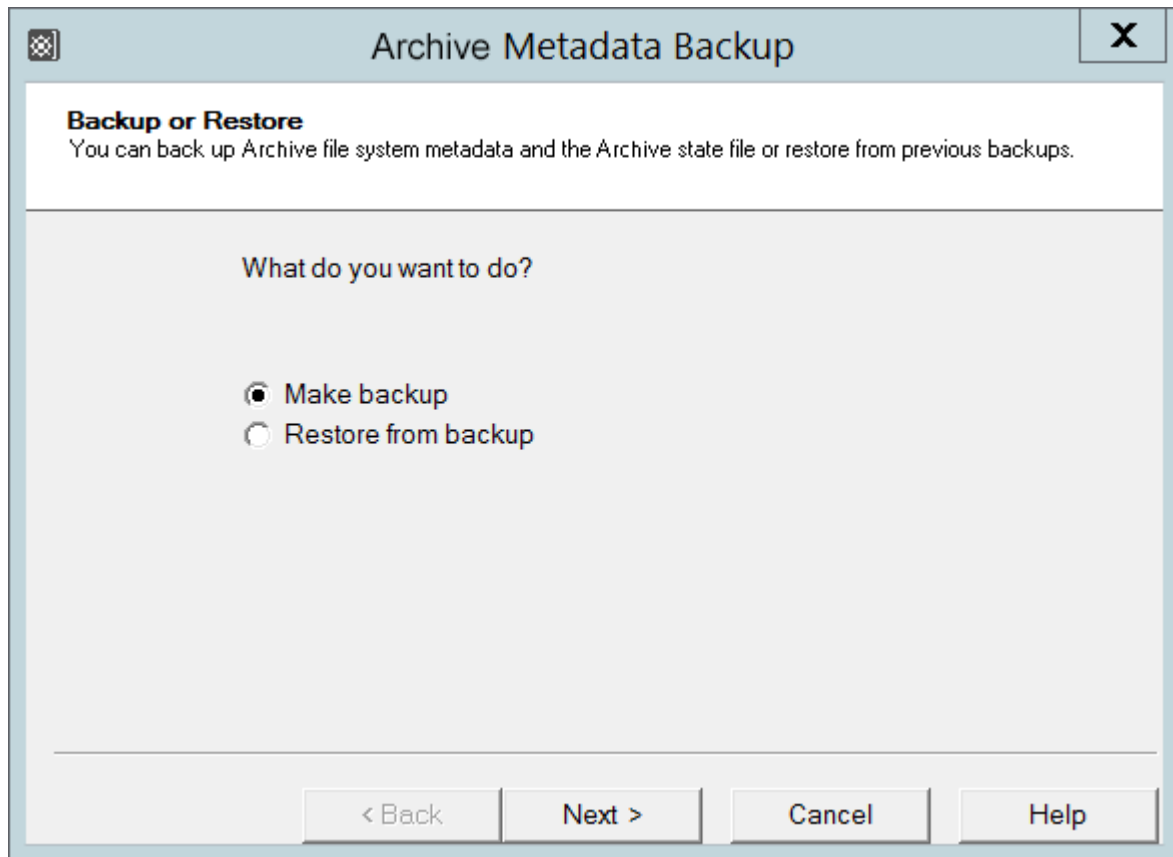
1. Click the Windows Start icon at the lower left hand corner of the screen.
2. Type "Metadata Backup".
3. Click the **Metadata Backup** entry in the list to the right of the screen.

6.2 Selecting Backup or Restore

The Metadata Backup program performs two types of operation:

- **Make backup** - makes a backup of the metadata in the system in its current state. See 'Making a Predefined backup' or 'Making a Custom Backup' below.
- **Restore from backup** - restores information from a backup file onto the cache disk. See 'Restore from backup' below.

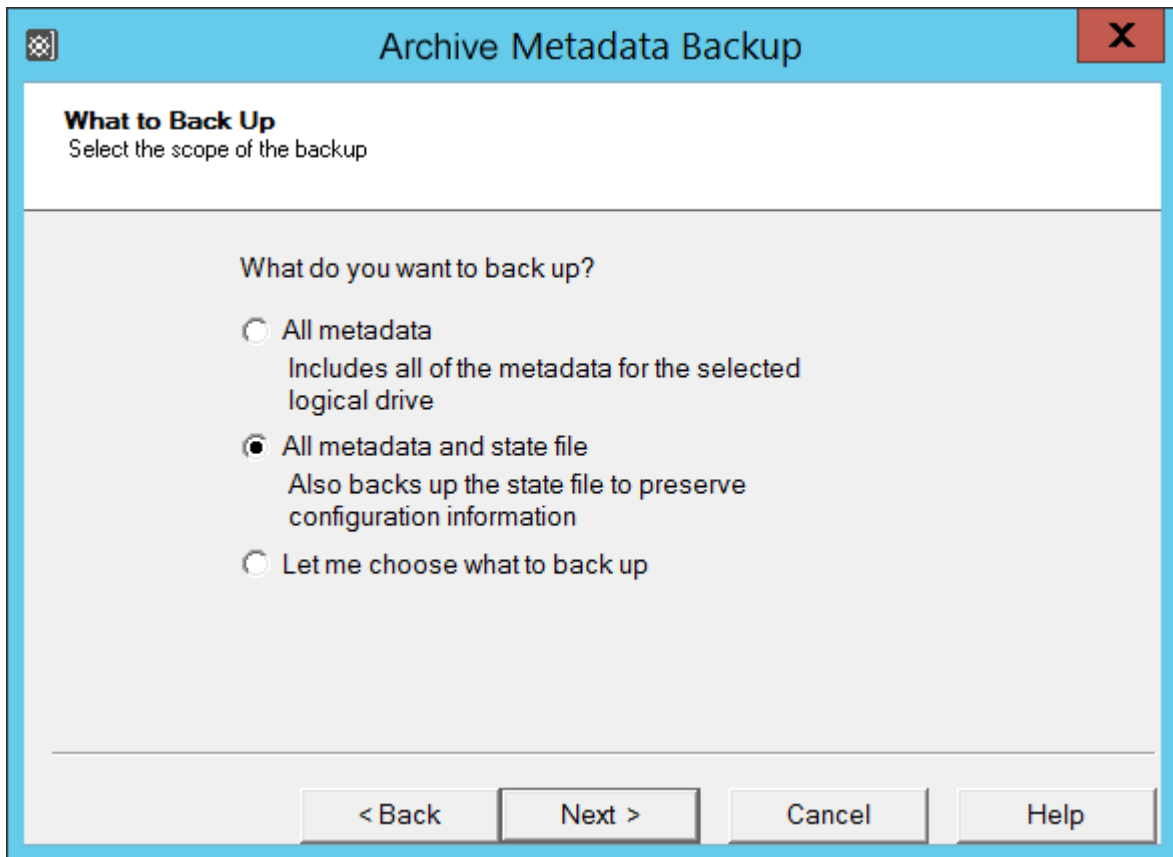
Select the desired option and click Next to continue.



6.3 Making a Predefined Backup

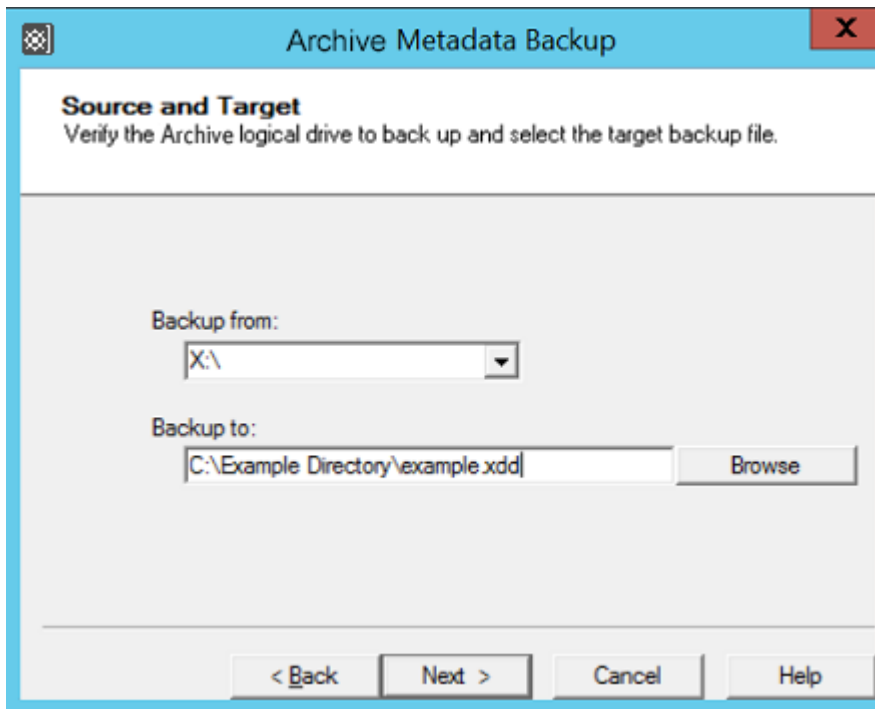
The instructions in this section describe how to perform a backup using one of the two predefined backup types. The section [Making a Custom Backup](#) describes how to use the **Let me choose what to back up** option to take more control over the backup. For example, a folder that is only used for temporary files may be excluded from the backup if the files it contains will not be required in future.

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

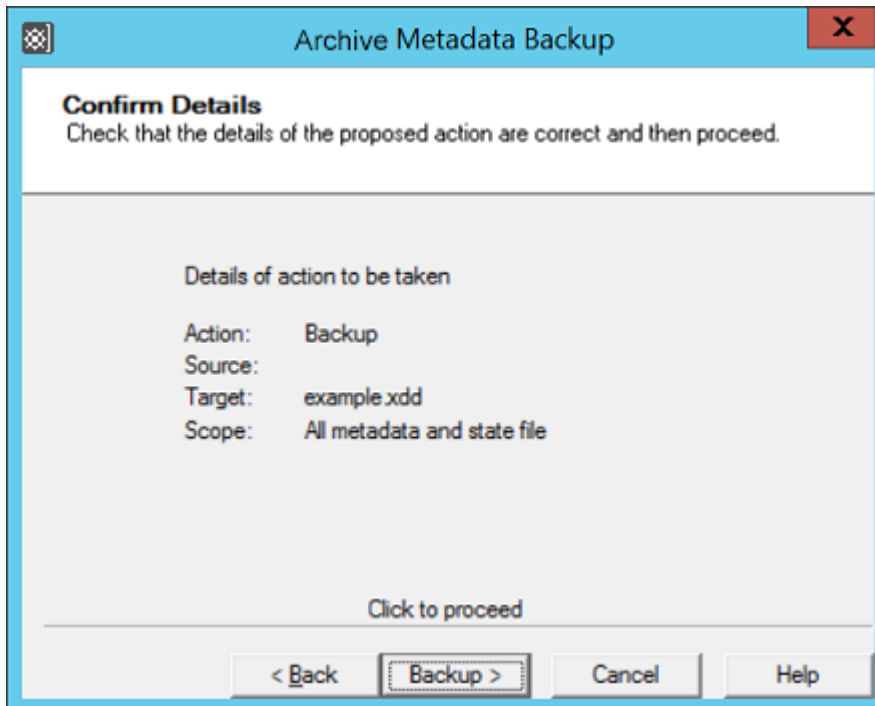


There are two predefined backup types. All metadata will back up all of the file system metadata, and All metadata and Archive state file will also include the State File.

1. Select **All metadata** or **All metadata and State File** as appropriate.
2. Click **Next** to continue.



1. Verify the cache disk drive to be backed up. Although XenData6 Server only supports one cache disk per system, you should verify that it has been correctly detected as the source for the backup.
 - 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 file name.
2. Click **Next** to continue.




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

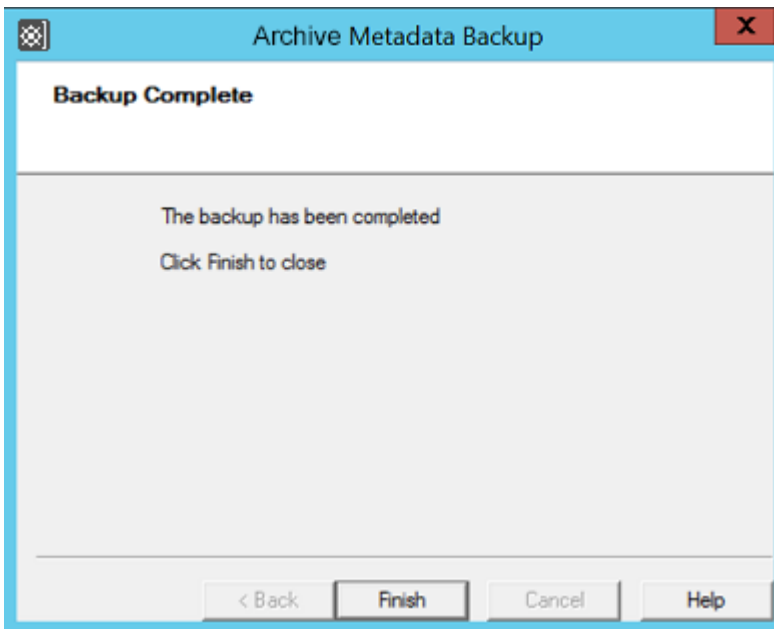
necessary.

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

A progress dialog box 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.

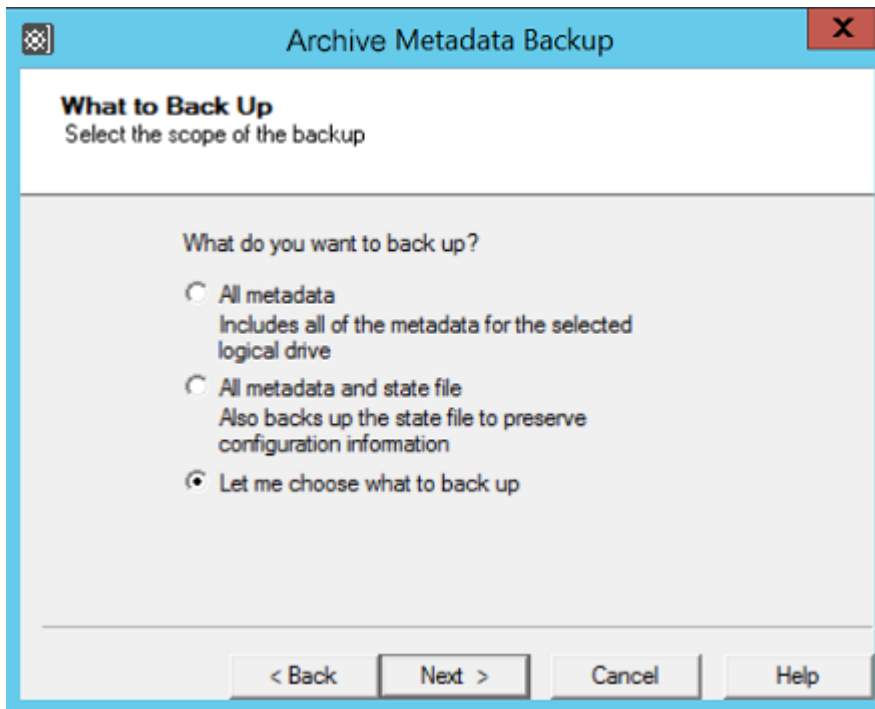


6.4 Making a Custom Backup

The instructions in this section describe how to perform a partial metadata backup, selecting what is included in the backup.

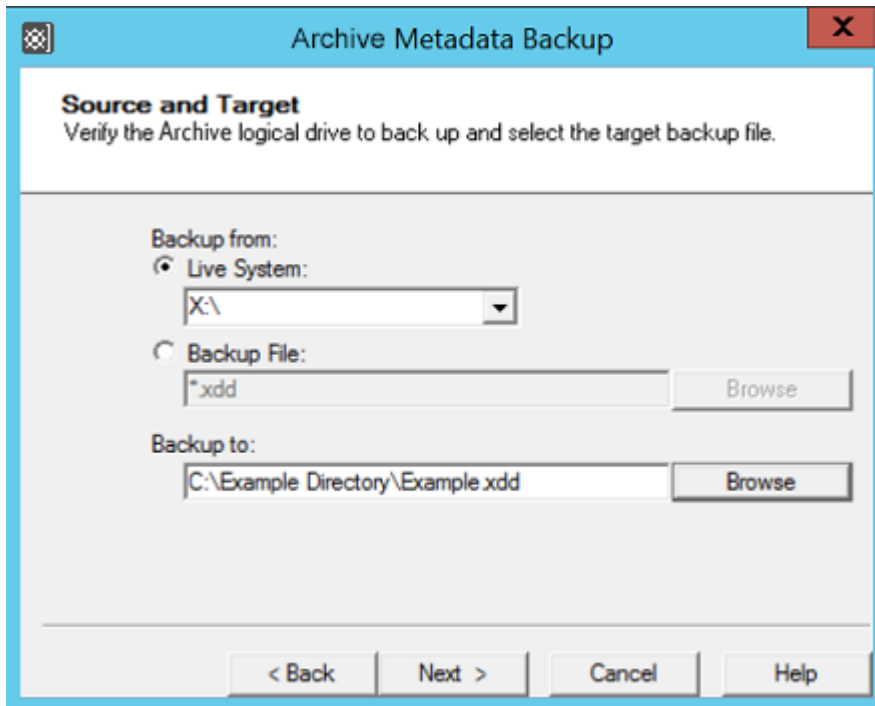
For example, a folder only used for temporary files may be excluded from the backup as the files it contains will not be needed following a system restore. It is also possible to create a sub-backup. This refers to creating a new backup file from an existing backup where the new backup contains only selected folders from the original backup file.

- Start the Metadata Backup program, select **Make backup** and click Next.



- The option **Let me choose what to back up** provides control over which file system metadata is backed up, and whether the State File is also included.

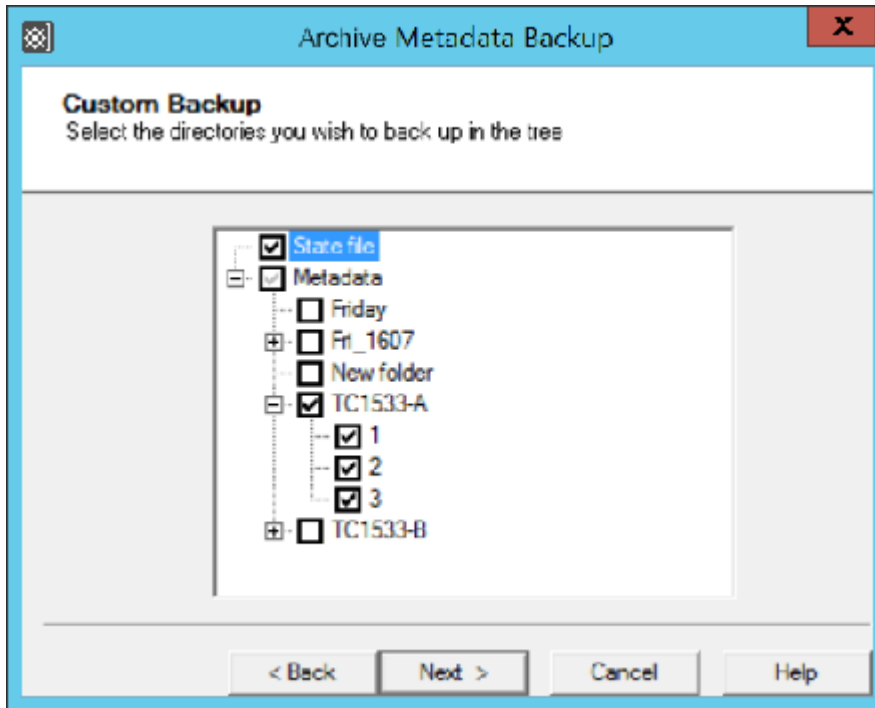
1. Select **Let me choose what to back up**.
2. Click **Next** to continue.



If a sub-backup of an existing backup file is being made, an existing backup file should be selected as

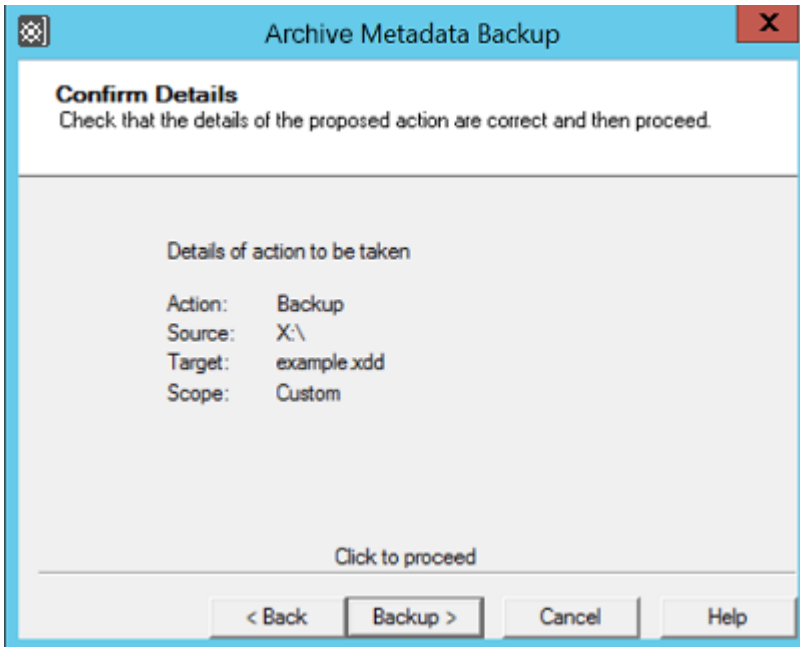
the source (the same file cannot be used as the target backup file). The Browse buttons can be used to assist in specifying the file.

1. Select **Live System** or **Backup File** as appropriate.
2. Either verify the cache drive letter 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 folder 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 sub-folder tree, and clicking on a "-" sign collapses it.


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 box appears that shows the backup progress, as illustrated below.



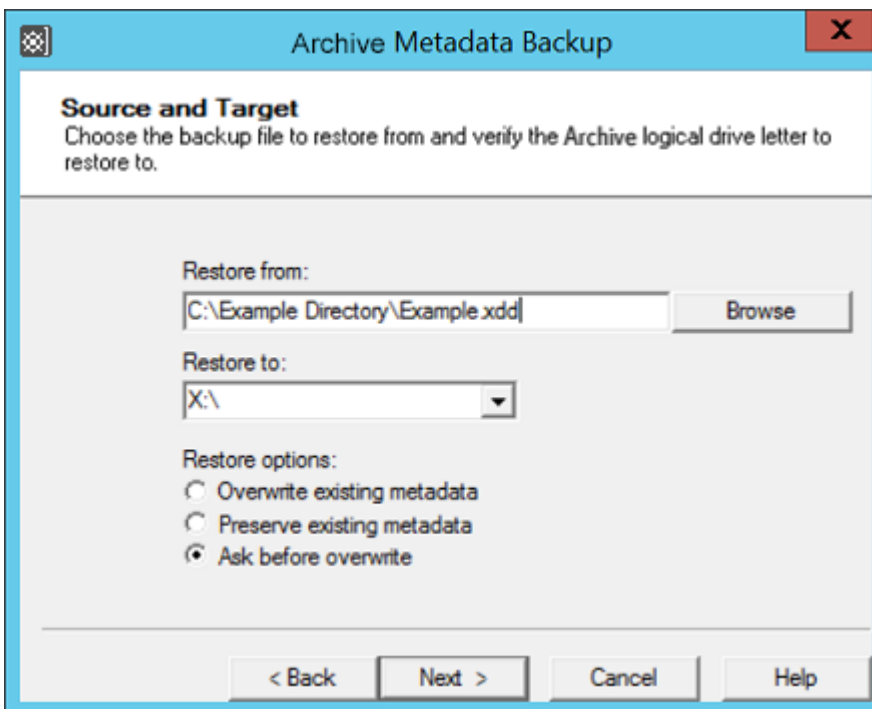
If the backup completed successfully, you will be presented with a confirmation page saying Backup Complete. Click **Finish** to dismiss the dialog box and exit the program.



6.5 Restoring from a 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 State File.

Either start the Metadata Backup program, select **Restore from backup** and click **Next** on the starting page, or double click on an Archive backup file (*.xdd) to display the Restore from backup prompt.

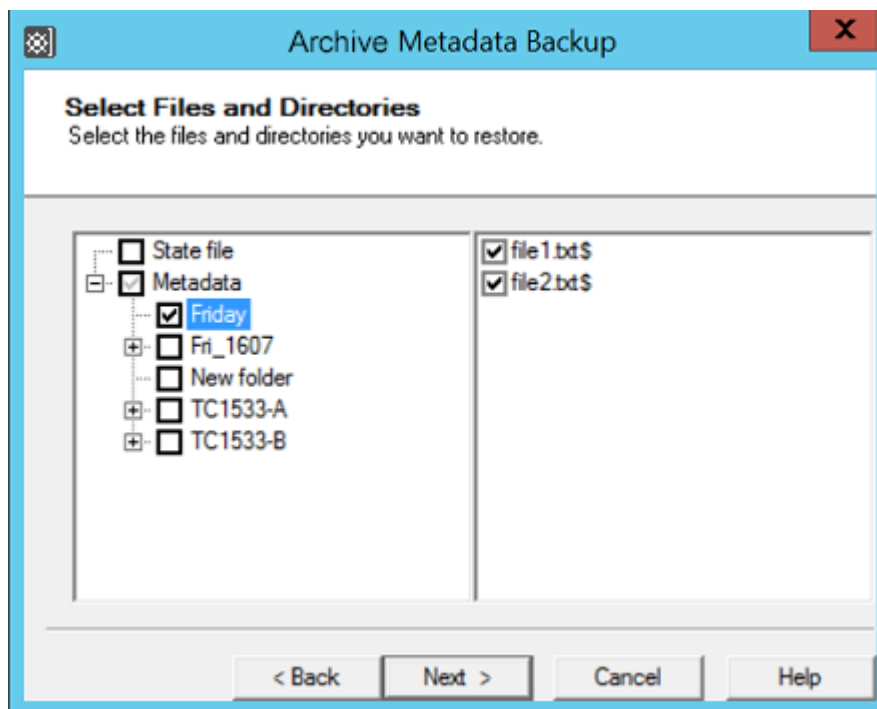


There are three restore options:

- **Overwrite existing metadata** - always writes metadata from the backup onto the cache disk, overwriting any metadata that is already present.
- **Preserve existing metadata** - will only write metadata for a particular file onto the cache disk if no metadata for that file is already present.
- **Ask before overwrite** - asks whether to overwrite existing metadata for each file whose metadata already exists, providing options to overwrite all of a certain category (for example, overwrite metadata where the existing metadata on the cache disk is currently invalid).

Although XenData6 Server only supports one cache disk per system, you should verify that it has been correctly detected as the target for the restore operation. Use the file dialog box opened by clicking **Browse** to browse for the source backup file.

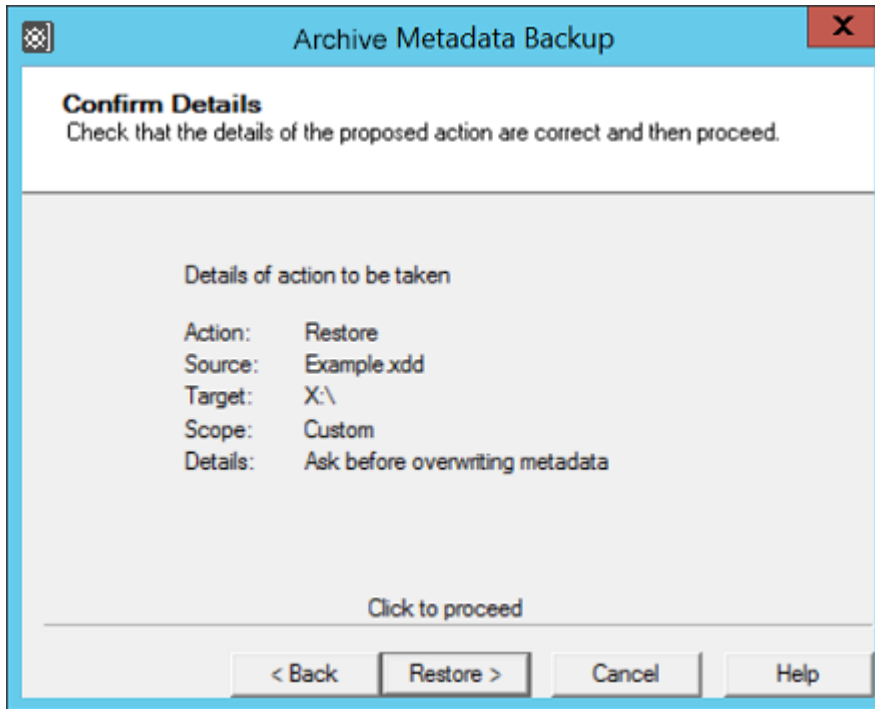
1. Specify the input backup file to restore from, or verify that the correct file name has been determined automatically.
2. Verify the cache logical drive to restore to.
3. Select the desired restore option.
4. Click **Next** to continue.



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 folder which needs to be traversed to reach checked items, but which will not itself be included is marked with a 'grayed out' check mark. When a folder is selected, the files within it are all selected by default, unless manually deselected.

Clicking on the "+" sign expands a sub-folder tree, and clicking on a "-" sign collapses it.


1. Select and deselect folders 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.

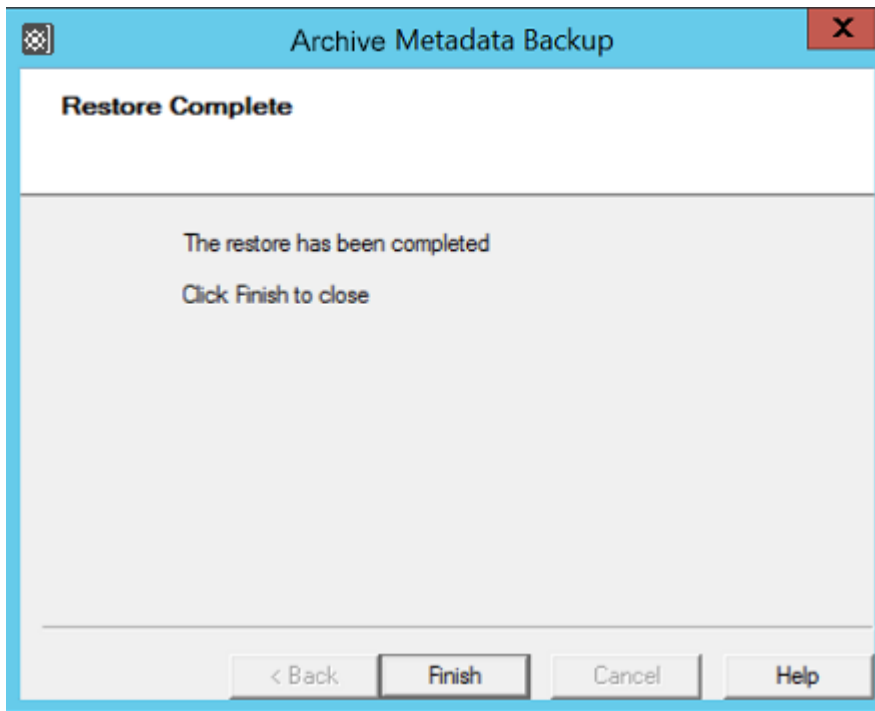
A progress dialog box will appear so that one can check the status of the restore operation. If the option to **Ask before overwrite** was selected during restore configuration, dialog boxes similar to the one shown below might appear, asking if existing metadata should be overwritten, and giving a category 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 dialog boxes 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 identical to that in the backup file, no overwrite dialog box will be displayed, no change is necessary and the file will be silently skipped.

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



Scheduler

7 Scheduler

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

- [Metadata Backup](#) which allows scheduling of full metadata backups including backup of the State File. It does not support scheduling of custom backups.
- [Deferred Writing](#) defers the initial writing of files to the primary replica of a Volume and allows the administrator to specify a scheduled time period when data can be archived. It is useful for prioritizing file restore operations during times of peak demand. **Note** that the scheduler does not define the timing for writing to replica tape cartridges; this must be configured as described [here](#).

7.1 Starting the Scheduler

For Windows Server 2008 R2 and Windows 7:

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

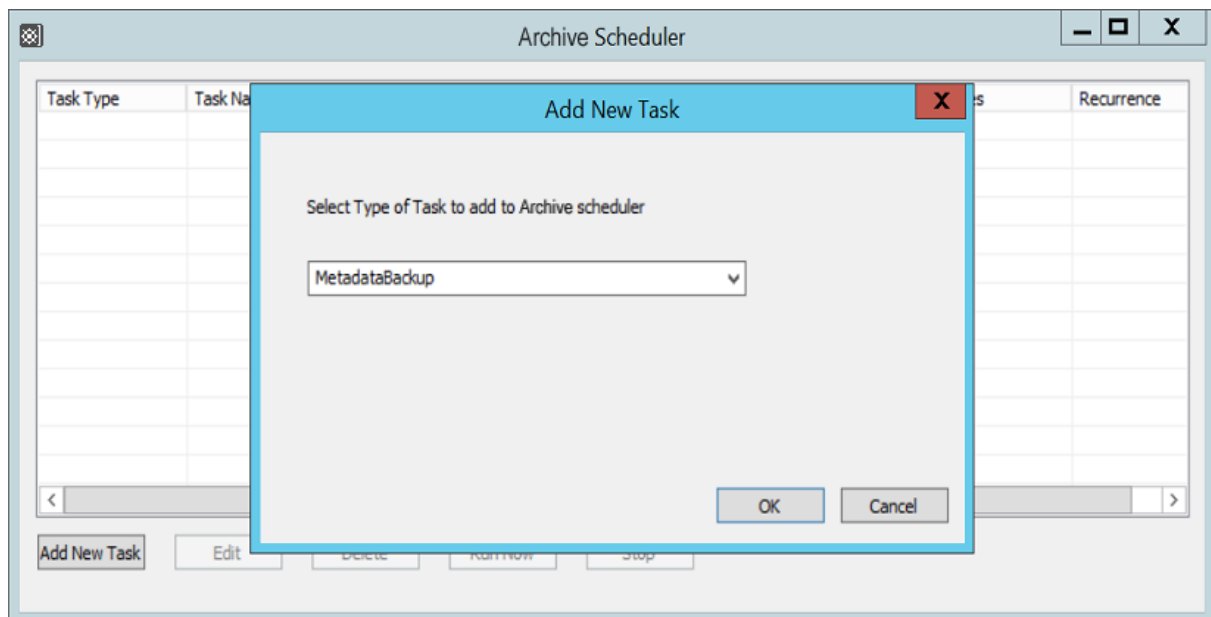
For Windows Server 2012R2:

1. Click the Windows Start icon at the lower left hand corner of the screen.
2. Type "Archive Scheduler".
3. Click the **Archive Scheduler** entry in the list to the right of the screen.

7.2 Adding a Task

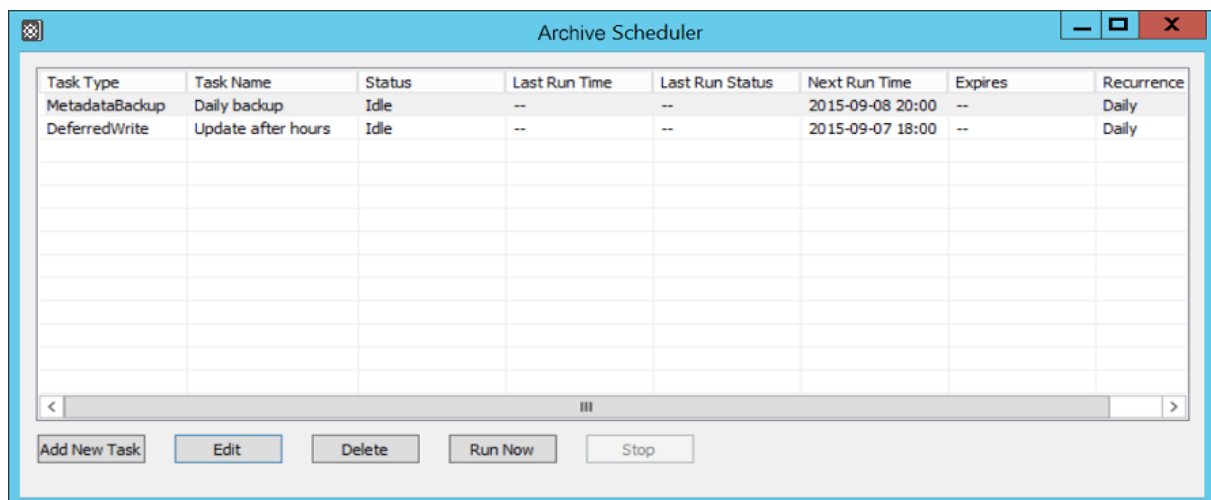
To Add a Task:

1. Start the Archive Scheduler
2. Click on **Add New Task** and then select the type of task from the drop down menu as shown below.



7.3 The Scheduler Status Display

An example of the Scheduler status display is shown below.



The display columns are as follows:

- **Task Type** - currently supported options are Metadata Backup and Deferred Write.
- **Task Name** - an optional parameter and can be left empty.
- **Status** - one of:
 - Idle – The task is not running. In this state an administrator can **Edit**, **Run Now** or **Delete** the task.
 - Running – The task is running and an administrator can **Stop** the task.
 - Locked – The task is being edited by another user. The task remains locked until the editing is complete.

- **Last Run Time** - shows the most recent date and time when the task was run. '--' indicates that the task has never run.
- **Last Run Status** - shows the result of the last task run. The status can be:
 - '--' – The task has never 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** - shows the date and time when the task will be run again. '--' indicates that the task will not be run again.
- **Expires** - optionally shows the date and time when a recurring task ends; '--' indicates that the task never expires.
- **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 Editing and Deleting Tasks

To Edit a Task

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

To Delete a Task

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

7.5 Starting and Stopping Tasks

In normal operation, the Archive Scheduler runs tasks automatically according to a predefined schedule. The [Scheduler Status Display](#) provide mechanisms to run a task "Now" and to stop a running task.

To Run a Task "Now"

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

To Stop a Running Task

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

Note that if a Metadata Backup Task is stopped by using the Stop button its 'Last Run Status' is set to 'FAIL' and no metadata backup file is created.

7.6 Scheduling Metadata Backup

Options for the Metadata Backup task are as follows:

- **Recurrence** is one of
 - 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.
- **Start** - sets the date and time for the first run of the task and defines the time and day of the week or date of the month when recurrence occurs
- **Expire** - optionally sets the date and time recurrence ends; '--' indicates that the task never expires.
- **Task Name** - is an optional parameter and may be left empty.
- **Chose directory path for backup** - determines where the backups will be located; the backup file name will be 'YYYYMMDDHHMM.xdd'. **Note** that the metadata backup task runs under the log-in ID used by the Archive Scheduler service (usually the Local System account). Ensure that the path entered here is accessible to that log-in ID (for example, the Local System account may not have access to network shares).
- **Delete previous backups** - removes previous backup files (with the extension XDD) upon successful completion of a metadata backup.

7.7 Scheduling Deferred Write

Deferred Write Task X

Recurrence

None
 Daily
 Weekly
 Monthly

Start

2015-09-07 13:32
 Expire
 2015-09-08 13:32

Task Name

My daily deferred volume write

Stop task if it runs longer than

30 minutes

Number of drives to use for deferred writes

1 drive

Enabling Deferred Write for a Volume Set will delay writing to primary replica until the Volume Set is updated using a scheduled task.
 Note that changing the deferred write status of a Volume Set from enabled to non-enabled will cause an immediate update.

Volume Sets with Deferred Write Enabled

Volume Set Identity	Volume Set Name
<input checked="" type="checkbox"/> 51B5E464-00000000	51B5E464-00000000

<--- --->

Volume Sets with Deferred Write Disabled

Volume Set Identity	Volume Set Name

Save Cancel

Options for the Deferred Write task are as follows:

- **Recurrence** is one of
 - 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.
- **Start** sets the date and time for the first run of the task and defines the time and day of the week or date of the month when recurrence occurs.
- **Expire** - optionally sets the date and time recurrence ends; '--' indicates that the task never expires.
- **Task Name** - is an optional parameter and may be left empty.
- **Stop task if it runs longer than** - defines the length of time the task can run.
- **Volume Sets with Deferred Write Enabled** - is a list of all the Volume Sets in the system that have deferred writing enabled. The Volume Sets that are selected with a check mark are controlled by this particular deferred write task. To completely disable deferred writing for a Volume Set, select it in the list and then click the '--->' button. This will trigger an immediate update of all deferred writes for the Volume Set.
- **Volume Sets with Deferred Write Disabled** - To enable deferred writing for a Volume Set, select it in the list and then click the '<---' button.

System Recovery

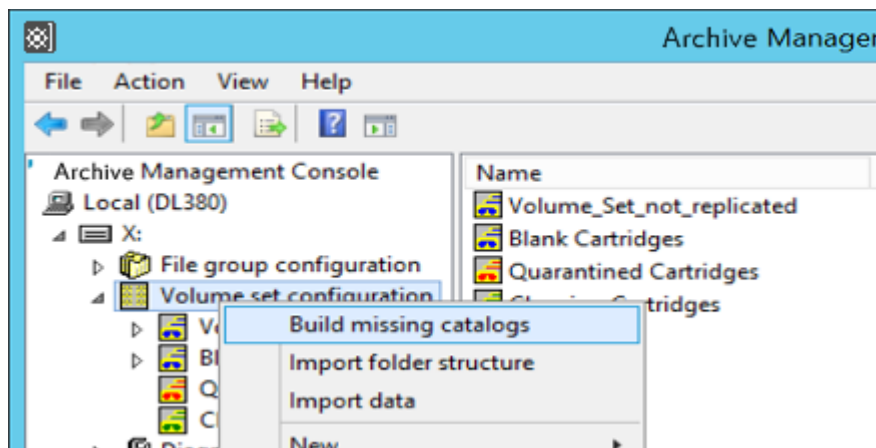
8 System Recovery

The preferred way to provide protection against hardware failure or other disaster is to use the [Metadata Backup](#) utility to take a periodic snapshot of the state of the system. The Metadata Backup utility generates a relatively small file that can rapidly restore a system's previous state (it can also be used to transfer system metadata from one XenData6 Server system to another). However, the data cartridge formats that are supported by XenData6 Server software are all completely self-describing and there is no need for any kind of external database or metadata to read the data stored on a cartridge. This property means that it is possible to rebuild a XenData6 Server system from the data cartridges alone, with no other form of backup. It is also possible to restore a system from an old metadata backup and then update the system from the cartridges that were written after the backup was made.

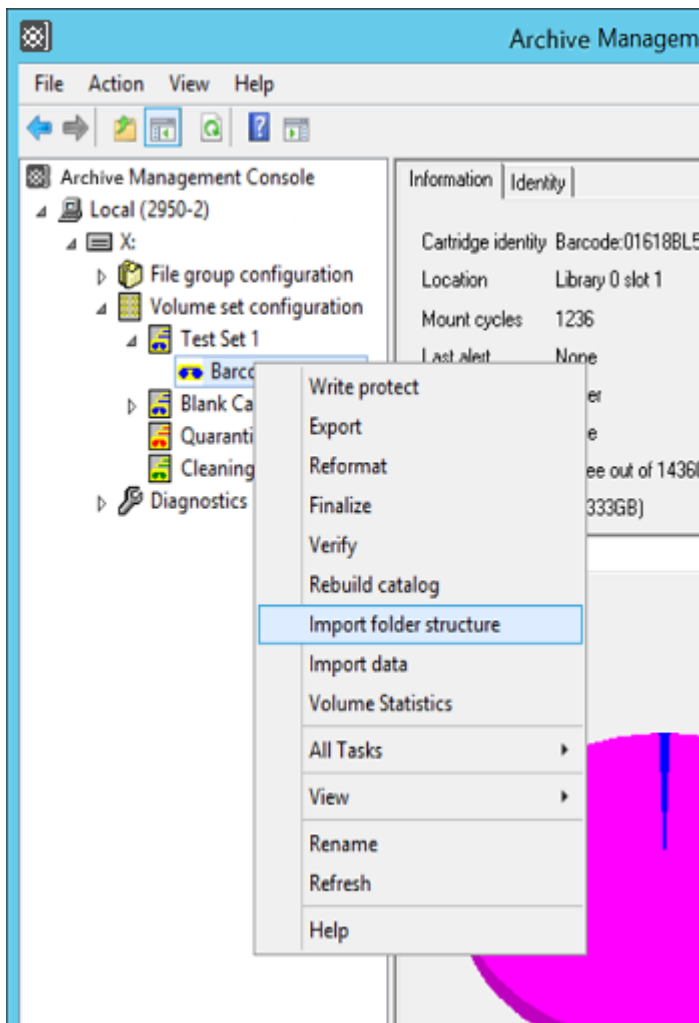
8.1 Rebuilding a System from Data Cartridges

Best practice dictates that the Metadata Backup utility should be used to protect an archive against catastrophic failure of the server or other disaster. However, all the data cartridge formats supported by XenData6 Server software are fully self-describing and in the event of there being no (or only a partial) metadata backup available, it is possible to rebuild the system from the data cartridges. This procedure provides an outline of the process for rebuilding a system from data cartridges.

1. If you have a metadata backup, use the [Metadata Restore](#) utility to load the available information into the new system.
 - o A full metadata backup includes the system configuration information (State File) at the time of the Metadata Backup. This includes information about blank cartridges. Cartridges that were blank at the time of the Metadata Backup may have been added to Volumes after the backup was made. Therefore you should remove information in the Metadata Backup about blank cartridges by [Forgetting](#) cartridges in the Blank Cartridge Set. This allows the system's automatic cartridge recognition algorithm to determine if cartridges are still blank or if they have been used.
2. Import the data cartridges into the system.
 - o Unknown cartridges will be recognized by the system and will show in the Archive Management Console. Cartridges written in the LTF5 or ODA formats and Finalized TAR cartridges will have Volume Contents Catalogs created on the cache disk.
3. Build the Volume Contents Catalogs for any non-Finalized TAR format tapes using the **Build Missing Catalogs** function shown below. This operation may take several hours because the entire cartridge must be scanned.



4. Use the **Import Folder Structure** or **Import Data** operations to publish files that were not present in a metadata backup to the file system interface. Files will be restored in the same state they are written on the cartridge. For example, files that were deleted will be created as deleted files visible with the [History Explorer](#).
 - Import Folder Structure loads file and folder information (metadata) into the system making the entire folder tree visible to users, but it does not restore the actual file data to disk. This operation is usually faster than Import Data.
 - Import Data loads file and folder metadata, but in addition, it selectively loads file data onto the cache disk, in accordance with the Disk Retention Rules for written files described in [Selecting Storage Options for a File Group](#). Files imported this way can be read directly from the disk cache without further access to a data cartridge.



8.2 In Case of Hardware Failure

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

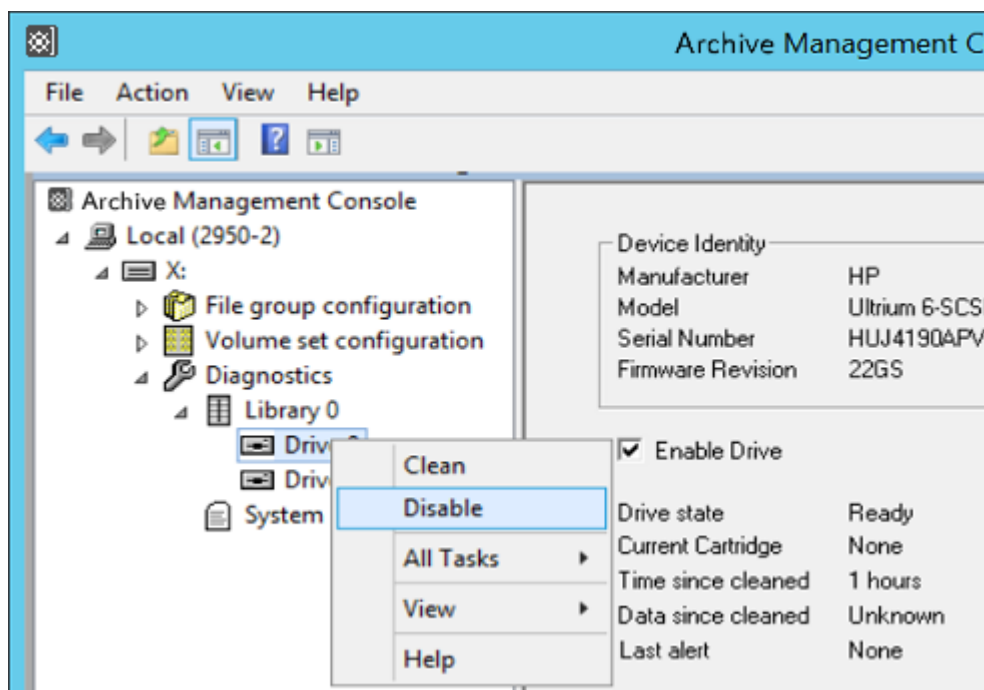
8.2.1 Options in Case of Library or Drive Failure

In normal operation, XenData6 Server software writes files to data cartridge Volumes as they are written to the cache disk. If this is not possible because of a hardware failure, the system will prevent further files from being written. If this behavior is undesirable (perhaps because there is no other space available for the data) then [Pending Write Mode](#) can be used to temporarily write data to the system cache disk.

8.2.2 Temporarily Disabling Hardware

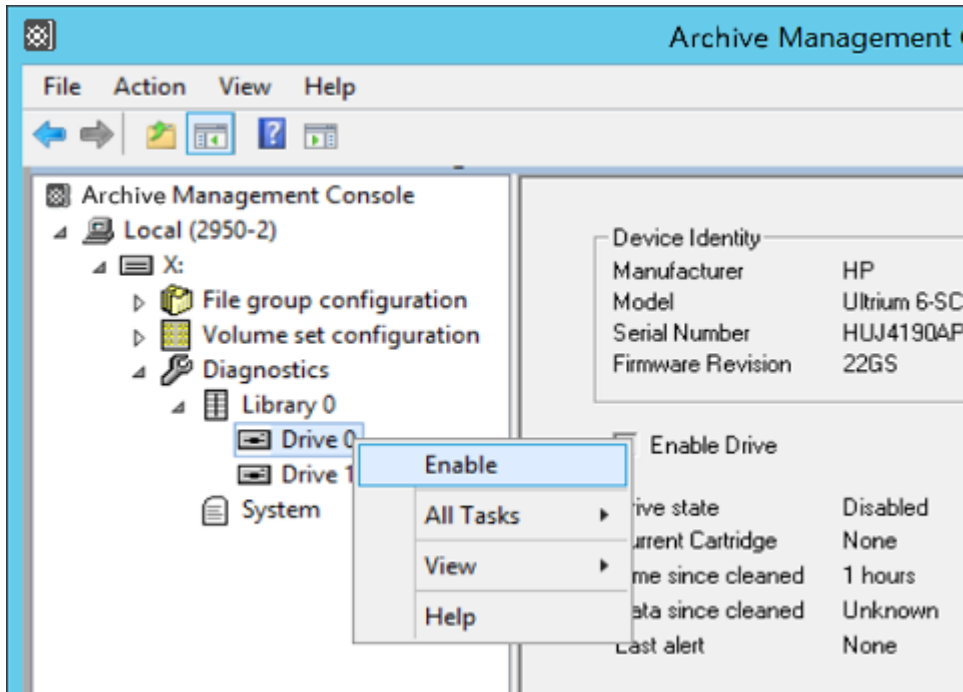
To Disable a Drive or Library:

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



To Re-enable a Drive or Library

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



Diagnostics & Maintenance

9 Diagnostics & Maintenance

XenData6 Server software maintains a comprehensive set of diagnostic information about the hardware under its control and allows the system administrator to selectively disable hardware that may need attention. Diagnostic information is logged in the Windows Event Log and the most recent event information can be found in the Diagnostics section of the Archive Management Console.

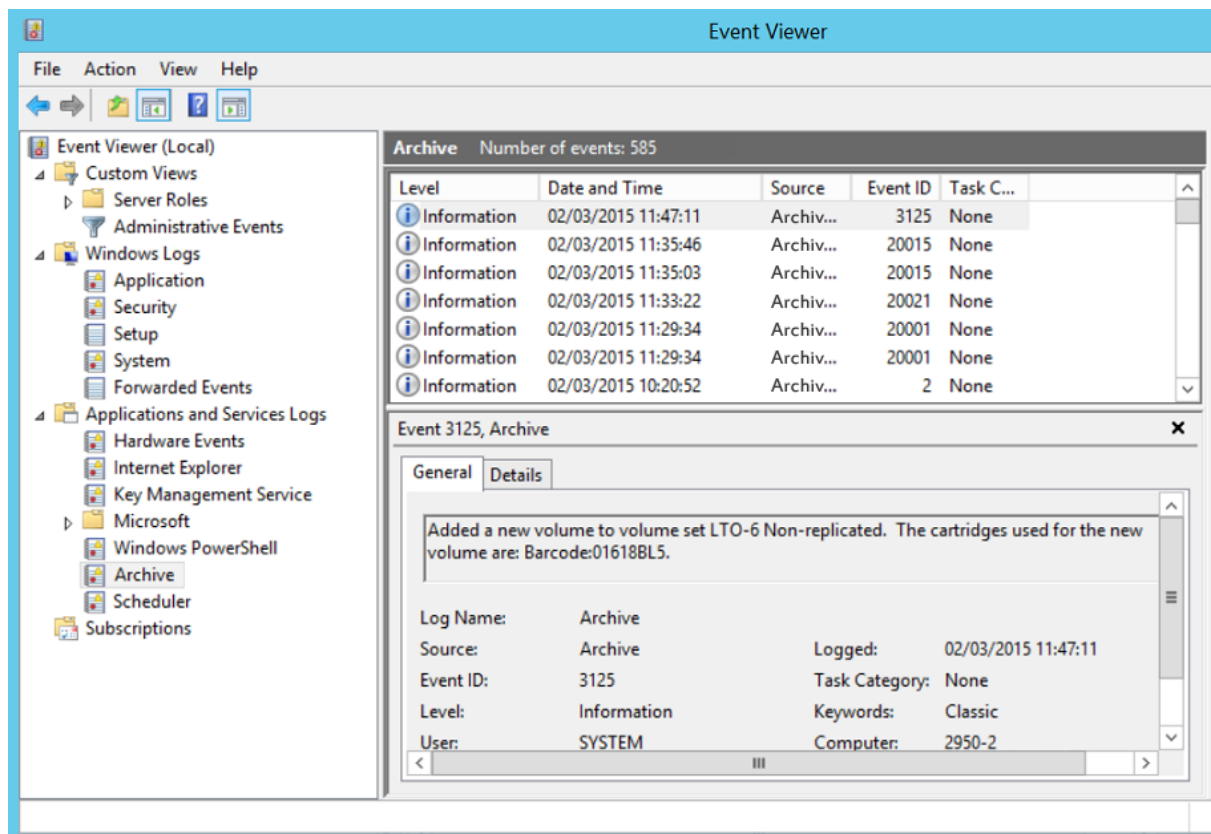
Tape drives employ a mechanism called Tape Alert to send maintenance information to external devices. XenData6 Server software monitors Tape Alert information and informs the user or takes other appropriate action (for example, cleaning a tape drive) where necessary.

9.1 Windows 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. The system also provides a comprehensive array of warnings for events such as running out of blank cartridges and it provides an audit of events like adding a new Volume to a Volume Set (as shown below). In general, if the system is not behaving as expected, the Windows Event Log is the first place that a system administrator should look.

To Open the Event Log:

1. Open the Windows Event Viewer.
2. Navigate to the XenData section of the Event Viewer as shown below.

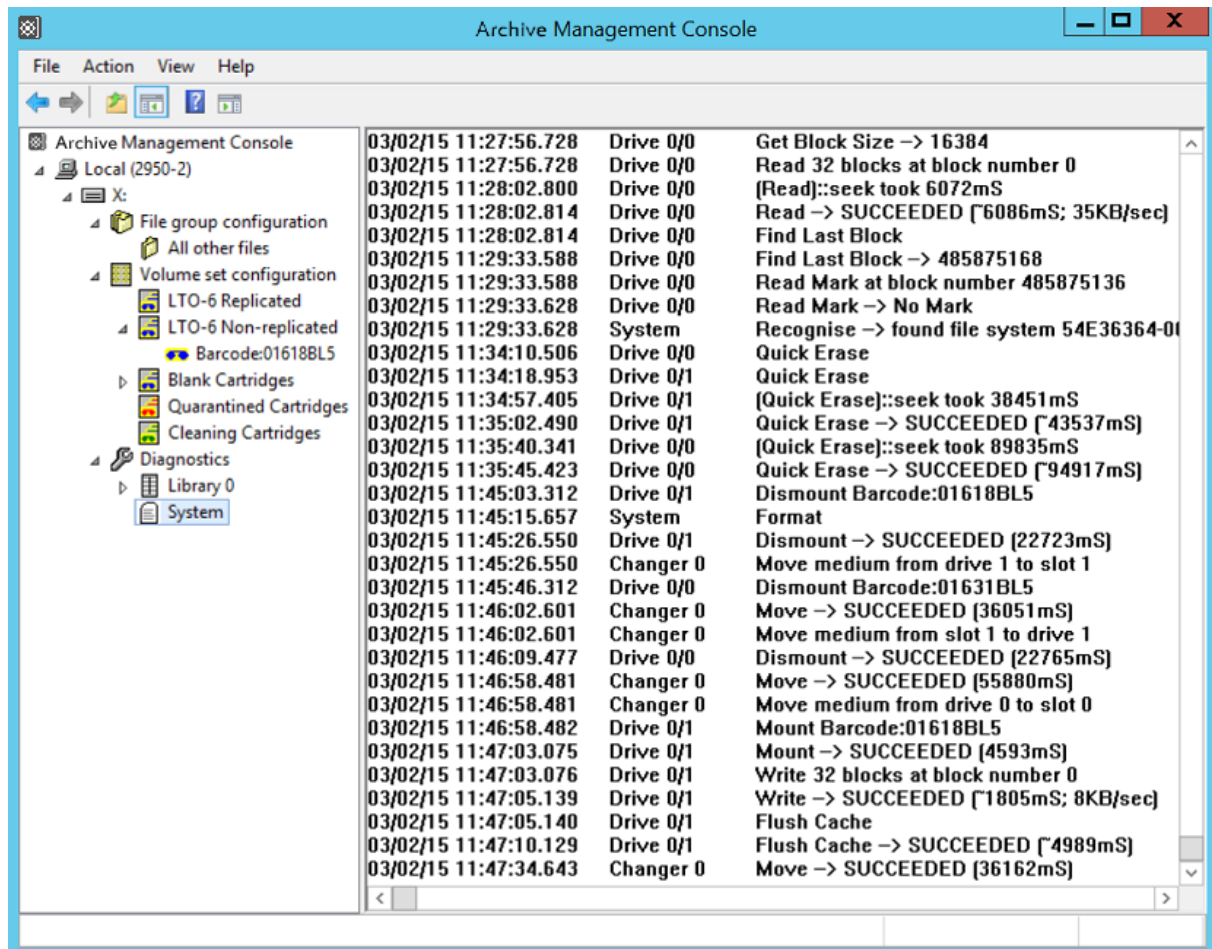


9.2 System Trace Log

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.

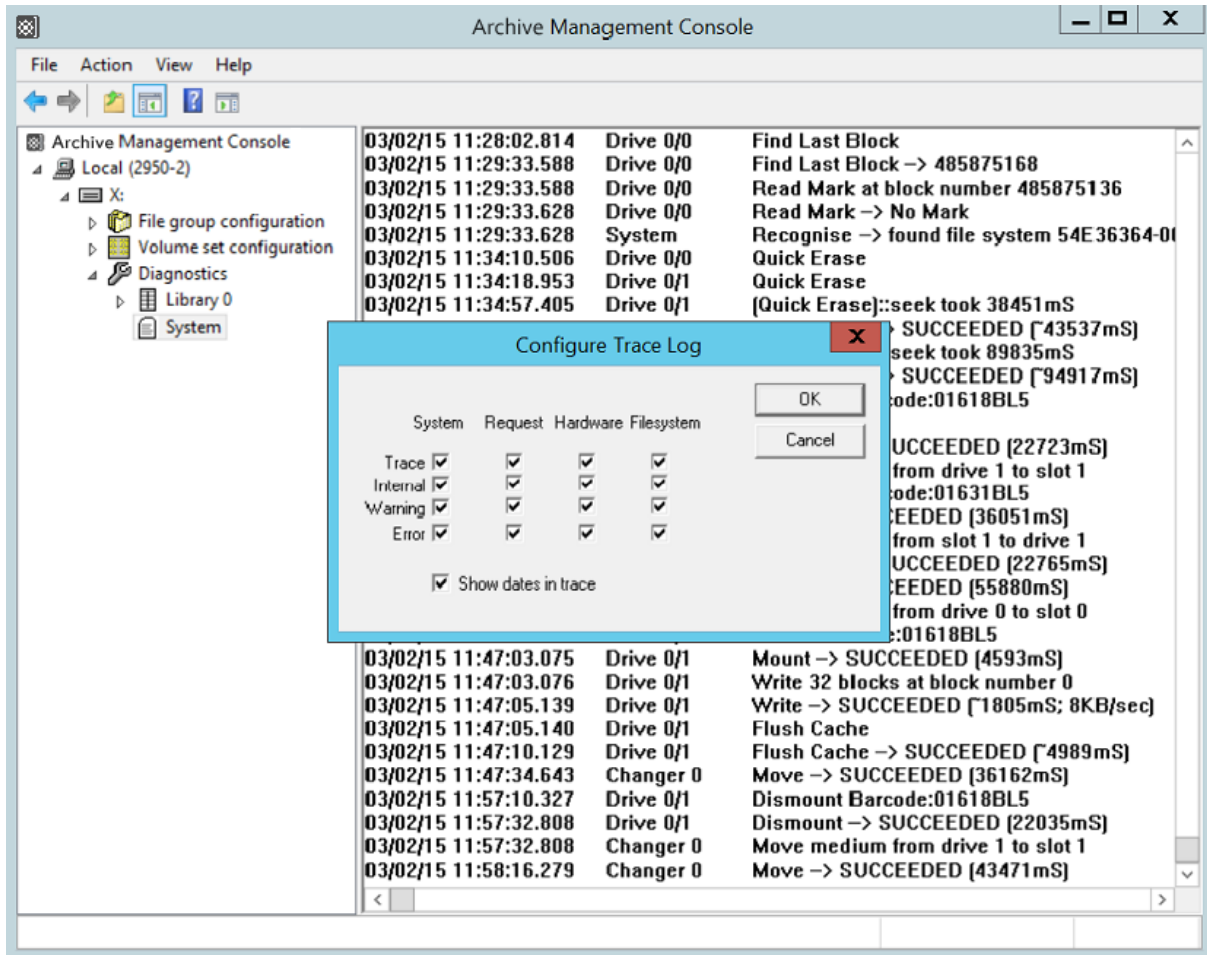
To Open the Trace Log

1. Open the Archive 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.



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.



Automatic Generation of Trace Files

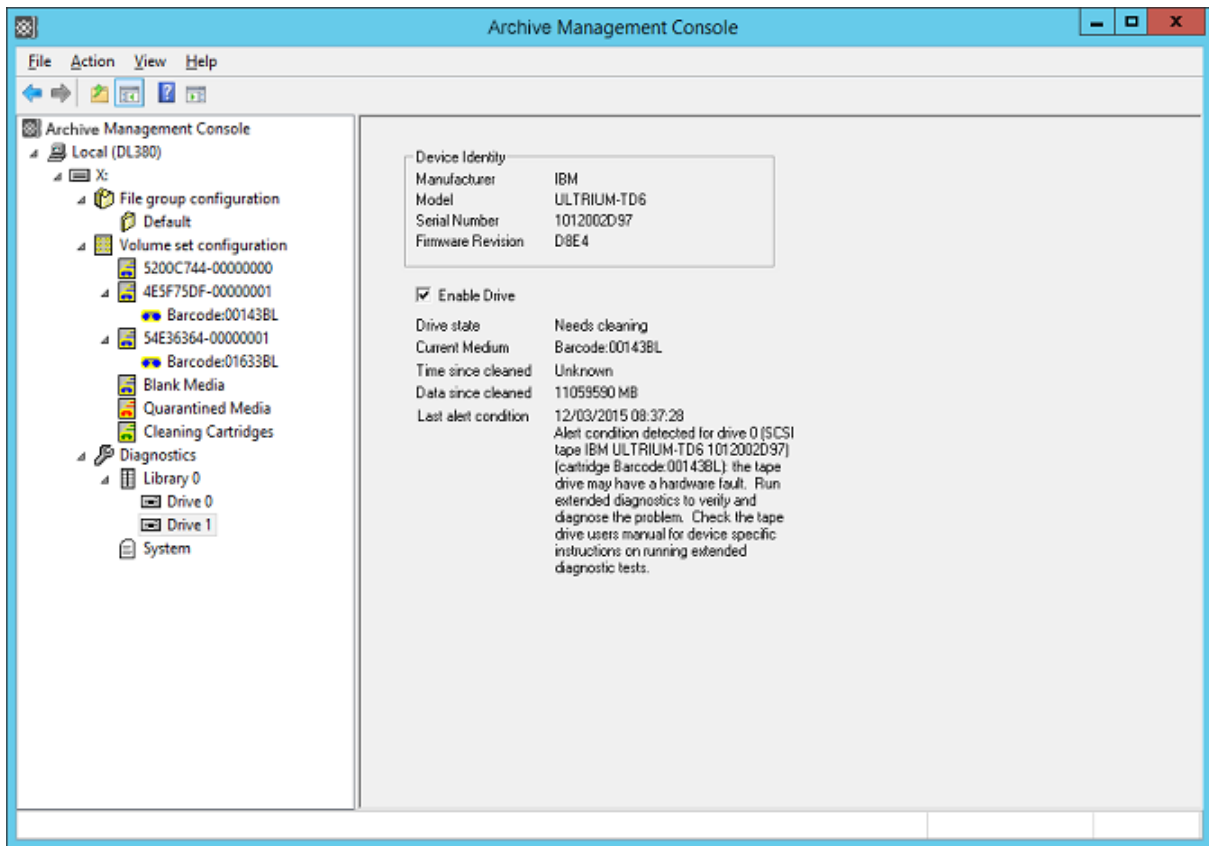
Whenever XenData6 Server software encounters an unexpected condition, it puts a message in the Windows [Event Log](#) and generates a trace file. The trace file contains a record of what the system was doing in the time, and is especially useful to assist support personnel in determining the cause of 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 Library and Drive Diagnostic Information

To Open the Diagnostics

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



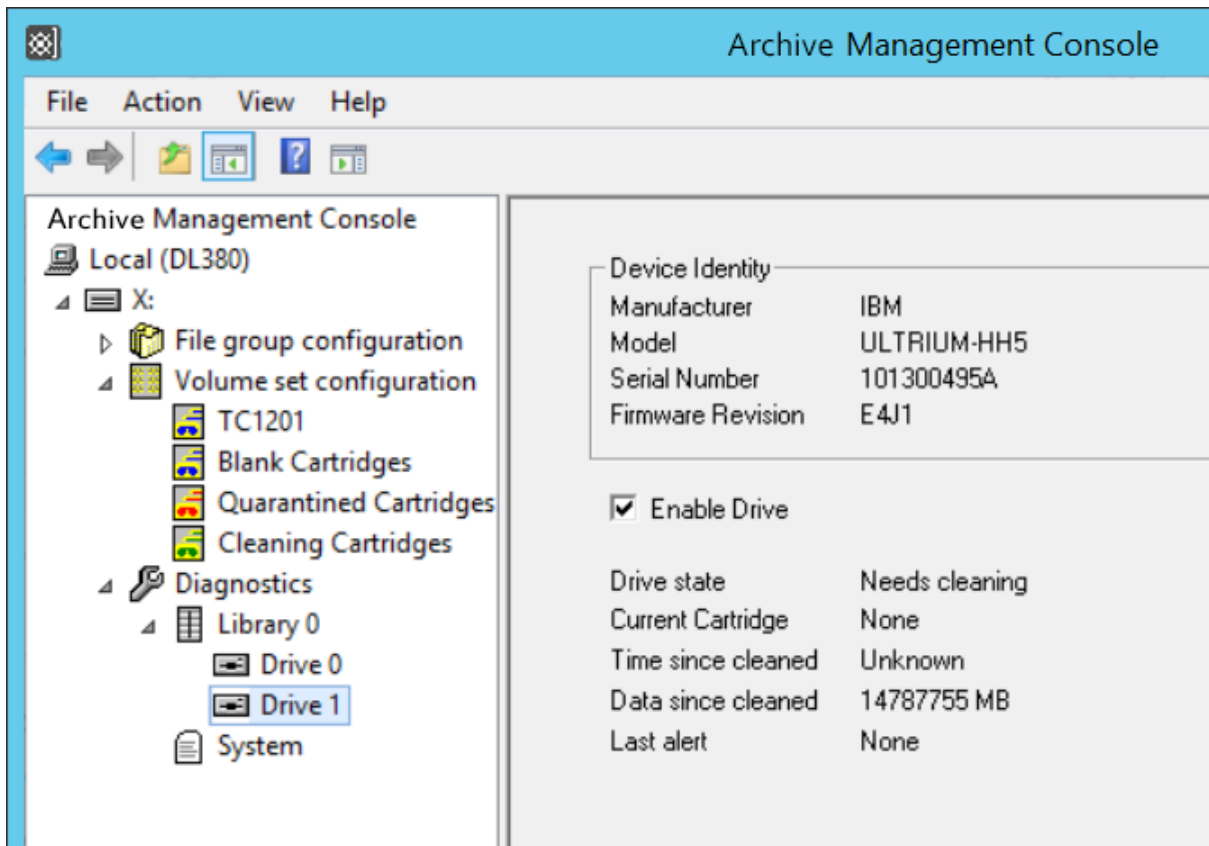
9.4 Cleaning Tape Drives

Most types of tape drive require periodic cleaning in order to maintain peak performance. Normally, a drive will issue a request for cleaning at the appropriate time and in the case of drives within robotic 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 library, the system will put a message in the Windows Event Log and in the Archive Management Console identifying that the drive requires cleaning.

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

To Manually Clean a Tape Drive within a Library

1. Open the Archive Management Console.
2. Navigate to the Diagnostics section.
3. Expand the Library icon and find the appropriate drive.
4. Right-click on the drive and select Clean.

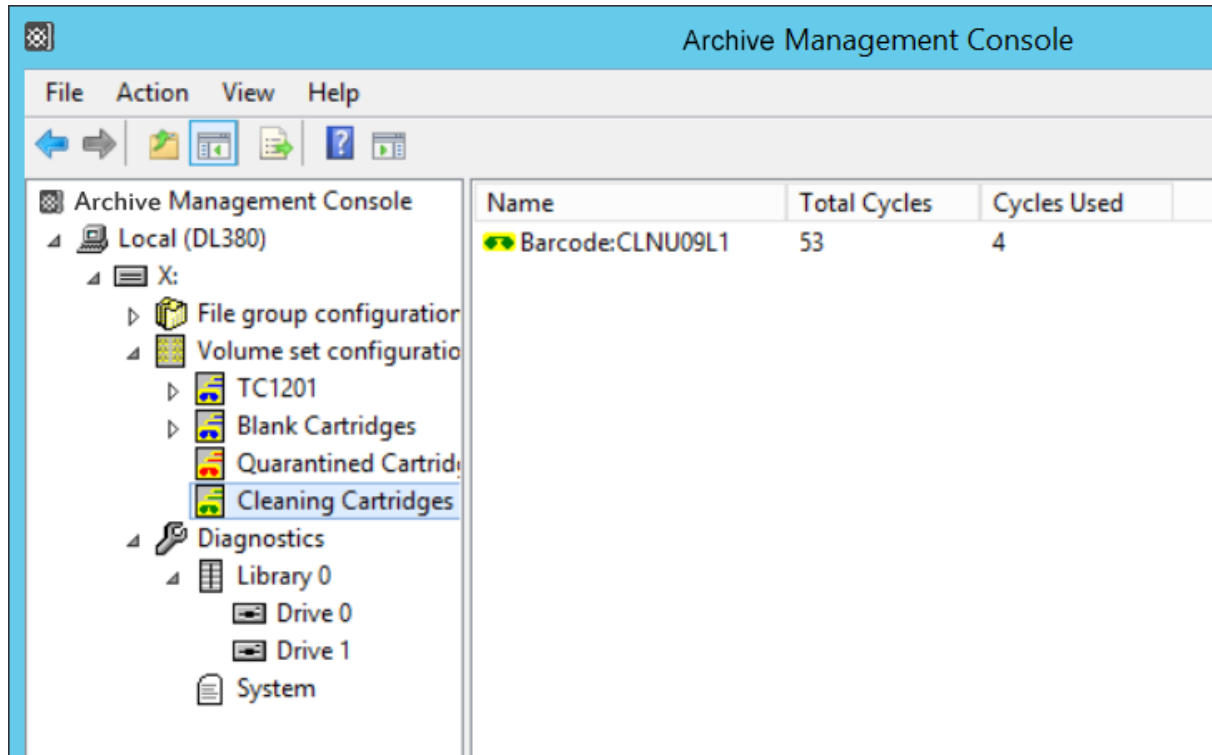


9.5 Displaying Information about Cleaning Cartridges

XenData6 Server software recognizes cleaning cartridges and displays information about them in the left pane of the Archive Management Console.

Cleaning cartridges are used to [clean tape drives](#) when necessary. The XenData6 Server system detects when drive cleaning is required and, in the case of tape drives within robotic libraries, automatically cleans the drive if a cleaning cartridge is available.

Note that cleaning cartridges have a specific barcode label format which always starts with CLN. This allows XenData6 Server software to recognize a cleaning cartridge without putting it in a drive and unnecessarily using a cleaning cycle.



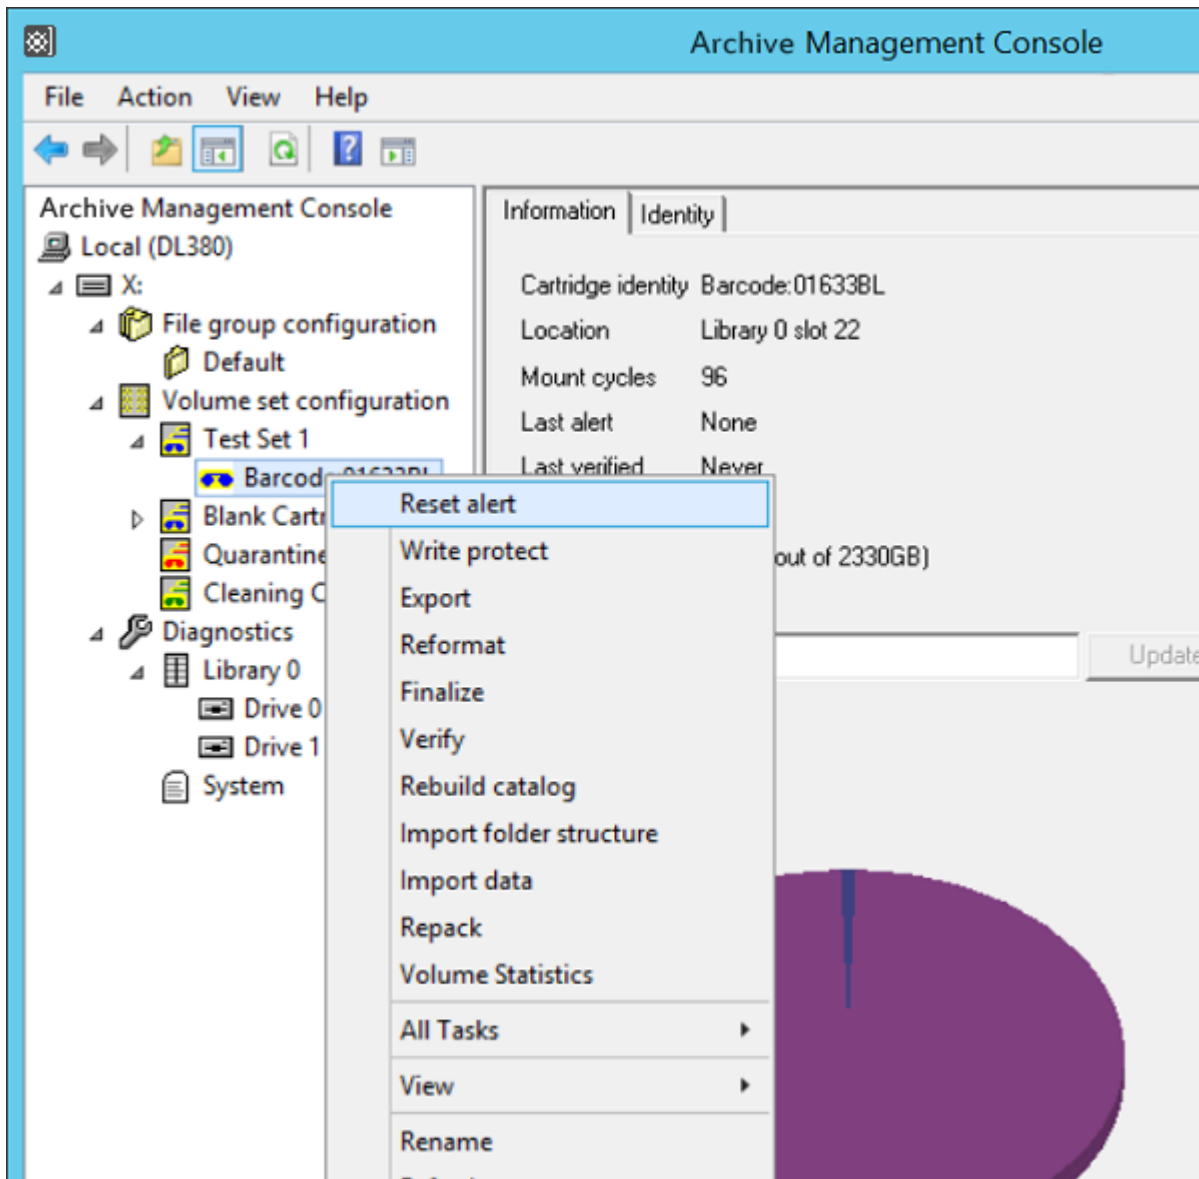
9.6 Cartridge Alert State

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

To Clear the Alert State for a Cartridge

1. Open the Archive Management Console.
2. Navigate to the affected cartridge or Volume.
3. Right click on the cartridge and select **Reset alert**.

Note that the Reset Alert option is only available for cartridges or Volumes 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 condition. In these cases, files written to the cartridge prior to the event will be readable. However, additional files cannot be written to the Volume and a new Volume must be [allocated](#) to the Volume Set to allow writing of additional files.

Reports

10 Reports

The Report Generator is included with XenData6 Server software. The available reports are:

- [File Search](#) - lists files that match a search term and identifies the data cartridges on which they are stored.
- [Recoverable Space](#) - lists the amount of space that can be recovered from a Volume using a Repack operation.
- [Data Cartridge Contents](#) - displays the files stored on a selected data cartridge.
- [UnArchived Files](#) - displays files that are not written to data cartridge Volumes but should be according to the current File Group rules.
- [Volume Contents](#) - 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.

10.1 Starting the Report Generator

For Windows Server 2008 R2 and Windows 7:

1. Click **Start**
2. Click **Programs** (or All Programs)
3. Click **XenData**
4. Click **Archive Report Generator**

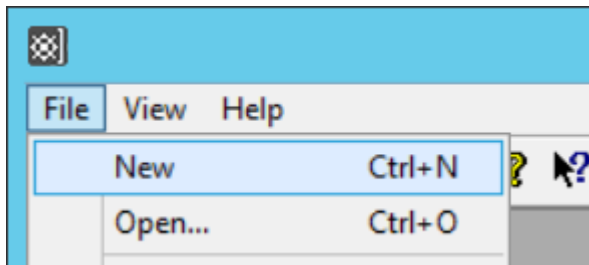
For Windows Server 2012R2:

1. Click the Windows Start icon at the lower left hand corner of the screen.
2. Type "Archive Report Generator".
3. Click the **Archive Report Generator** entry in the list to the right of the screen.

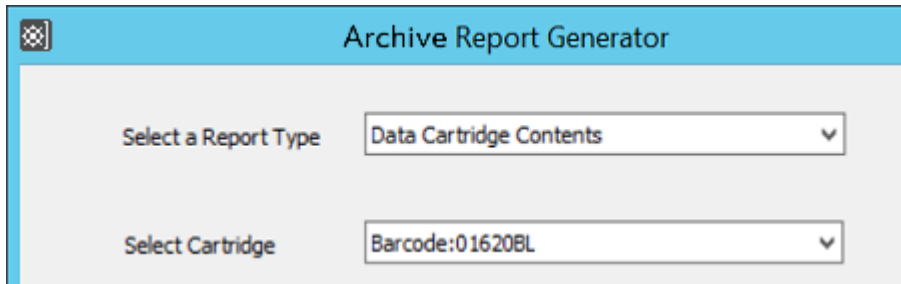
10.2 Creating, Saving and Restoring Reports

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.

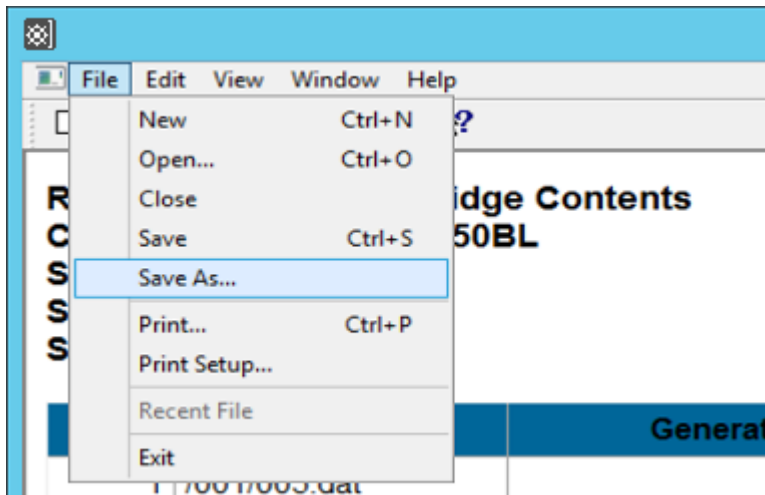


Please refer to the applicable section below for instructions on the selected report type.

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 or other applications.

To save a report, select the **File** and **Save As** menu options as shown below.

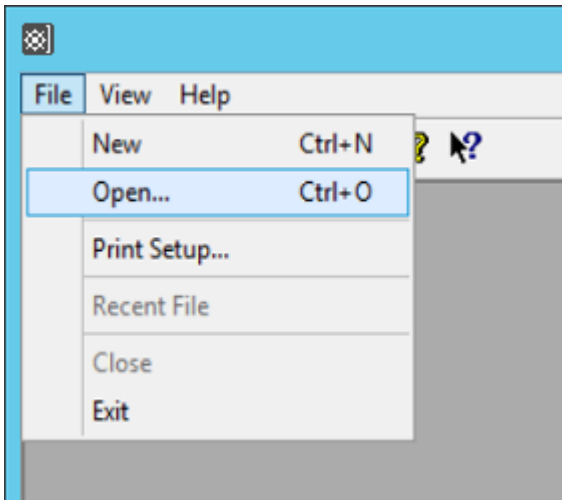


Then browse to the required location, select the file name and format and then click **Save**.

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 the **File** and **Open** menu options as shown below.

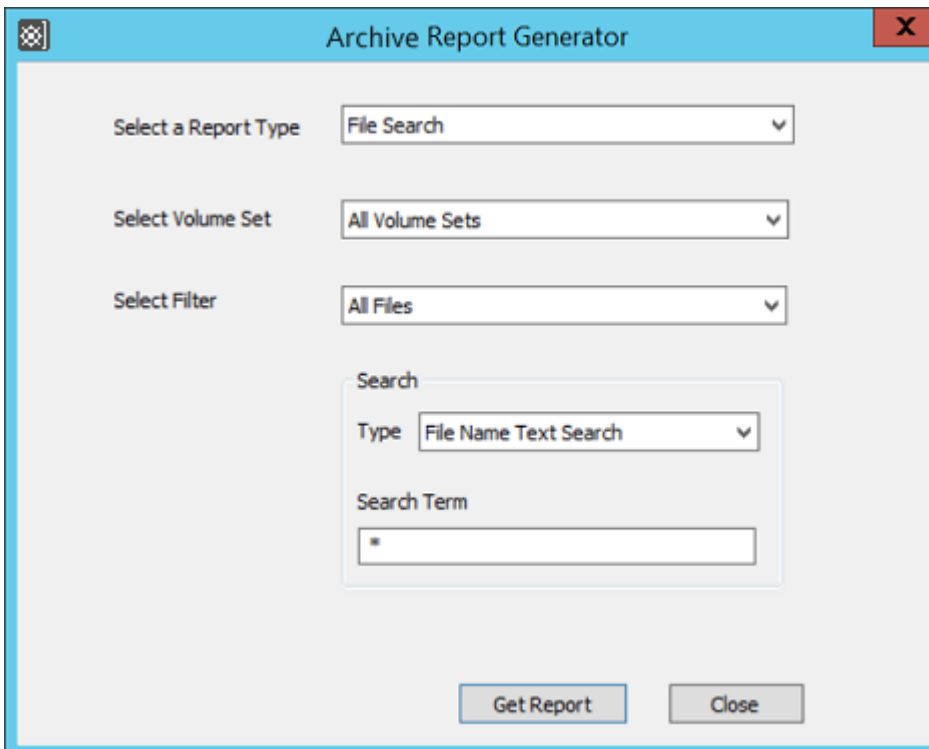


Then browse to the location of the saved report, select the required XRG file and open it.

10.3 File Search Report

To Run a File Search Report

1. Start the Report Generator.
2. Select the **File** and **New** menu options.
3. Select **File Search** as the report type.



The File Search Report lists archived files that match a search term and identifies the data cartridges

where they are stored. The search may be limited to a single Volume Set or may include all Volume Sets. The displayed report can be filtered in the following ways:

- **All Files** - displays all files including deleted files, old versions of files and renamed files.
- **Only Current Files** - displays only the files that can be accessed via the Windows file system interface and excludes deleted files, old versions of files and renamed files.
- **Only Deleted Files** - displays only deleted files.

It is possible to search using a simple text search or using a Regular Expression. When **File Name Text Search** is chosen, the search option supports wild cards.

Select the Volume Set, the filtering options and search type and term then click **Get Report**.

Note: A File Search Report will search only in the Volumes that have a Volume Contents Catalog file stored on the system cache disk.

10.3.1 Interpreting a File Search Report

An example of a File Search Report is shown below.

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

No	File Name	Generation	Version	File Size (bytes)	Barcode	Status
1	/Venice/002632433169_Venice Toma to Zaccaria_May_2006.avi	0	1	1,487,700,480	Barcode:00144BL5	Current
2	/Venice/002632433169_Venice Toma to Zaccaria_May_2006.mov	0	1	682,195,238	Barcode:00144BL5	Current
3	/Venice/002645458536_Venice taking the bus_May_2006.avi	0	1	989,532,160	Barcode:00144BL5	Current
4	/Venice/002645458536_Venice taking the bus_May_2006.mov	0	1	428,431,097	Barcode:00144BL5	Current
5	/Venice/002678933456_Venice Grand Canal_May_2006.avi	0	1	411,132,928	Barcode:00144BL5	Current
6	/Venice/002678933456_Venice Grand Canal_May_2006.mov	0	1	141,976,961	Barcode:00144BL5	Current

Done NUM

The File Search Report lists archived files that match a search term. The display columns are described below.

- **No** - the sequence number of the file in the display sorted by either date or file name, as defined by

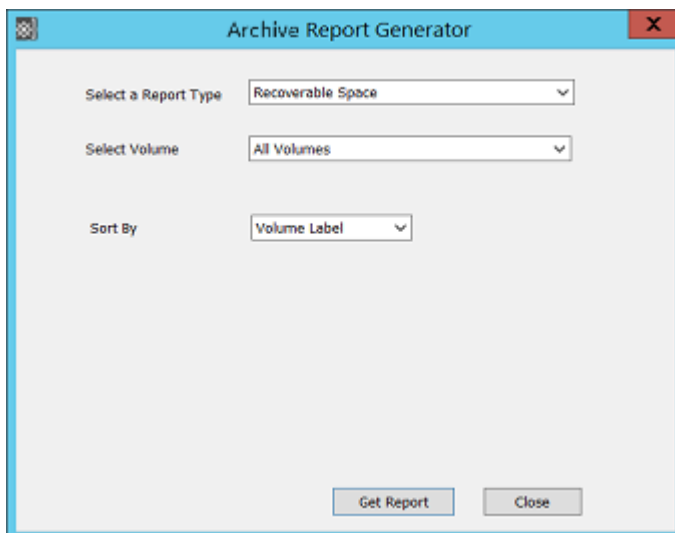
the **Sort by** selection.

- **File Name** - 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, the generation number is set to 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 data, XenData6 Server 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 Volume, this column displays the file size stored on the selected cartridge or Volume followed by the total size of the file in bytes.
- **Bar code** - this is the barcode label of the cartridge that contains the file.
- **Type** The status of the file is displayed as one of the following:
 - **Current** - this is the most recent version of the file, accessible through the archive drive letter.
 - **Renamed** - the file has been renamed and is now accessible under a different name.
 - **Deleted** - the file has been deleted and is no longer accessible except via the **History Explorer**
 - **Overwritten** - the file has been overwritten and this version is no longer accessible except via the **History Explorer**.

10.4 Recoverable Space Report

To Run a Recoverable Space Report

1. Start the Report Generator.
2. Select the **File** and **New** menu options.
3. Select **Recoverable Space** as the report type.



Select the Volume, the sort option then click **Get Report**.

10.4.1 Interpreting 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 the [Repack](#) operation which recovers tape space used by deleted files and old versions of files.

No	Volume Label	Barcode	Used Space (bytes)	Available Space (bytes)	Recoverable Space (bytes)
1	51B5E464-00000000-5600073B	Barcode:00144BL5	1,089,036,353,536	1,531,410,644,992	0
2	51B5E464-00000000-560125A8	Barcode:01639BL5	5,348,786,176	2,615,098,212,352	0

The display columns are described below.

- No - the sequence number of the Volume in the display.
- Volume Label - the Volume Label for the cartridge.
- Bar code - this is the barcode label of the cartridge.
- Used Space (bytes) - the total amount of space consumed on the Volume.
- Available Space (bytes) - the available free space for the Volume.
- Recoverable Space (bytes) - the amount of space recoverable by using a repack operation.

10.5 Data Cartridge Contents Report

To Run a Data Cartridge Contents Report

1. Start the Report Generator.
2. Select the **File** and **New** menu options.
3. Select **Data Cartridge Contents** as the report type.

Archive Report Generator

Select a Report Type: Data Cartridge Contents

Select Cartridge: Barcode:00144BL5

Select Filter: All Files

Search

Type: No Search

Sort by: Date

Search Term:

Get Report Close

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

- **All Files** - displays all files on the cartridge including deleted files, old versions of files and renamed files.
- **Only Current Files** - displays only the files that can be accessed via the Windows file system interface and excludes deleted files, old versions of files and renamed files.
- **Only Deleted Files** - displays only deleted files.

You can further filter the results by specifying a **Search Type (File Name Text Search or Regular Expression Search)** . When **File Name Text Search** is chosen, the search option supports wild cards.

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

Note: A Cartridge Contents Report will search only on Volumes that have a Volume Contents Catalog file cached on the system.

10.5.1 Interpreting a Cartridge Contents Report

An example of a Cartridge Contents Report is shown below.

Archive Report Generator - [Report4]

File Edit View Window Help

Report Type: Data Cartridge Contents
 Cartridge: Barcode:01639BL5
 Showing: All Files
 Search Type: None
 Sorted by: Date

No	File Name	Generation	Version	File Size (bytes)	Date Archived	Type
1	/Venice_2/002632433169_Venice Toma to Zaccaria_May_2006.avi	0	1	1,487,700,480	Sep 22 2015 11:01	Current
2	/Venice_2/002632433169_Venice Toma to Zaccaria_May_2006.mov	0	1	682,195,238	Sep 22 2015 11:01	Current
3	/Venice_2/002645458536_Venice taking the bus_May_2006.avi	0	1	989,532,160	Sep 22 2015 11:01	Current
4	/Venice_2/002645458536_Venice taking the bus_May_2006.mov	0	1	428,431,097	Sep 22 2015 11:01	Current
5	/Venice_2/002678933456_Venice Grand Canal_May_2006.avi	0	1	411,132,928	Sep 22 2015 11:01	Current
6	/Venice_2/002678933456_Venice Grand Canal_May_2006.mov	0	1	141,976,961	Sep 22 2015 11:01	Current

Done NUM

The display columns are described below.

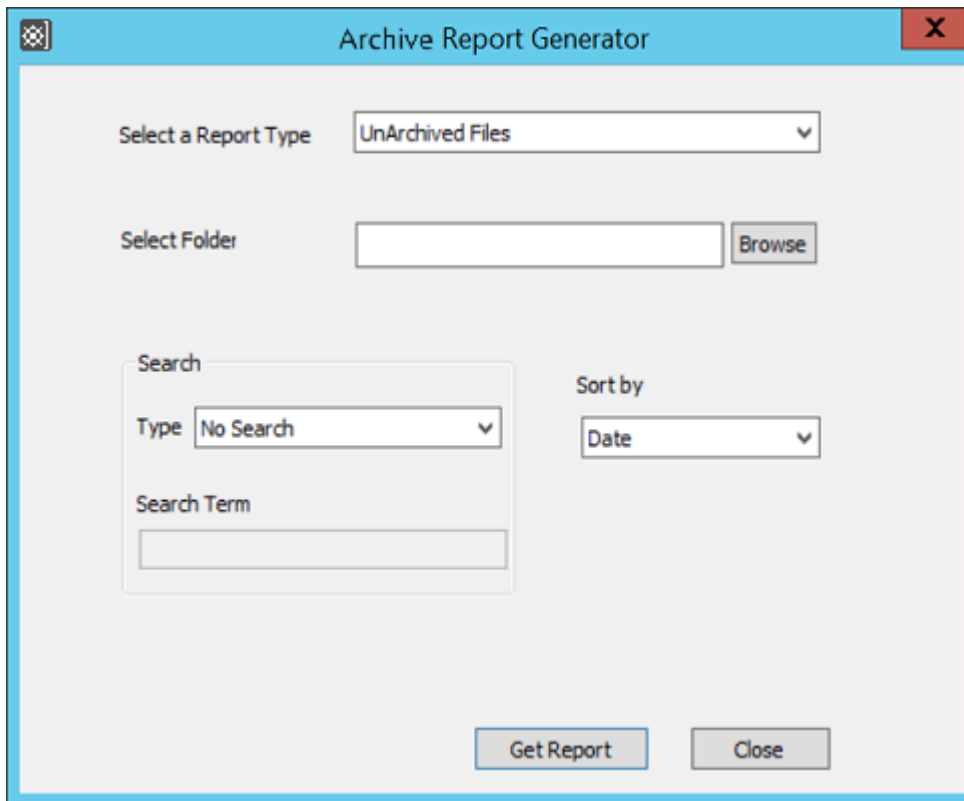
- **No** - the sequence number of the file in the display sorted by either date or file name, as defined by the **Sort by** selection.
- **File Name** - the file name including full path from the root of the archive logical drive letter.
- **Generation** - when a file of a given name and path is first created, the generation number is set to 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 data, XenData6 Server 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 Volume, this column displays the file size stored on the selected cartridge or Volume followed by the total size of the file in bytes.
- **Date Archived** - the date and time the file was archived.
- **Type** The status of the file is displayed as one of the following:
 - **Current** - this is the most recent version of the file, accessible through the archive drive letter.
 - **Renamed** - the file has been renamed and is now accessible under a different name.
 - **Deleted** - the file has been deleted and is no longer accessible except via the **History Explorer**.
 - **Overwritten** - the file has been overwritten and this version is no longer accessible except via the **History Explorer**.

10.6 UnArchived Files Report

The UnArchived Files Report lists files which are not fully archived to data cartridge Volumes and that should be archived according to the current File Group rules.

To Run an UnArchived Files Report:

1. Start the Report Generator.
2. Select the **File** and **New** menu options.
3. Select **UnArchived Files** as the report type.



Select a folder as the start point of the search (all sub-folders will be included in the search). You can further filter the results by specifying a **Search Type (File Name Text Search or Regular Expression Search)** which will filter the displayed results. When **File Name Text Search** is chosen, the search option supports wild cards.

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

10.6.1 Interpreting an UnArchived Files Report

An example of an UnArchived Files Report is shown below.

The screenshot shows a window titled 'Archive Report Generator - [Report1]'. The menu bar includes File, Edit, View, Window, and Help. Below the menu is a toolbar with various icons. The main content area displays the following report information:

Report Type: UnArchived Files Report
Search in Folder: X:
Search Type: None
Sorted by: Date

No	File Name	Generation	Version	Replica	Barcode	Status
1	/Test1/01.dat	0	1	1	Barcode:00144BL5	Archived
				2	Barcode:01639BL5	Not Archived
2	/Test1/02.dat	0	1	1	Barcode:00144BL5	Archived
				2	Barcode:01639BL5	Not Archived
3	/Test1/03.dat	0	1	1	Barcode:00144BL5	Archived
				2	Barcode:01639BL5	Not Archived
4	/Test1/04.dat	0	1	1	Barcode:00144BL5	Archived
				2	Barcode:01639BL5	Not Archived
5	/Test1/05.dat	0	1	1	Barcode:00144BL5	Archived
				2	Barcode:01639BL5	Not Archived

The status column shows 'Archived' for the first replica of each file and 'Not Archived' for the second replica. The status is split across two rows for each file entry. The status 'Archived' is also visible in the bottom right corner of the table area.

The display columns are described below.

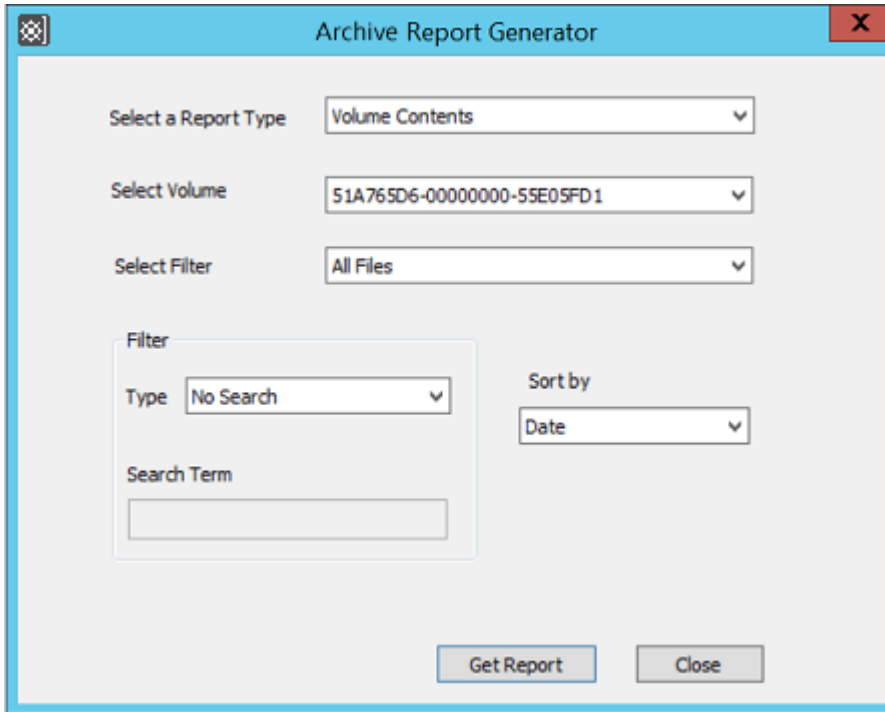
- **No** - the sequence number of the file in the display sorted by either date or file name, as defined by the **Sort by** selection.
- **File Name** - the file name including full path from the root of the archive drive letter.
- **Generation** - when a file of a given name and path is first created, the generation number is set to 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 data, XenData6 Server 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** - in the case of files where the File Groups rules are associated with replicated Volume Sets, this identifies the replica whose status is displayed in columns further to the right.
- **Bar code** - available in cases where a data cartridge has been assigned for the file, for example when a write operation started but did not complete.
- **Status** - a 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 data was written to a data cartridge but XenData6 Server software was unable to verify that the operation had completed successfully.
 - **Archived** - this instance of the file is archived correctly (replicas in this state are only displayed when at least one other instance of the file is not archived correctly to a replica cartridge).

10.7 Volume Contents Report

The Volume Contents Report lists the contents of the selected cartridge.

To Run a Volume Contents Report

1. Start the Report Generator.
2. Select the **File** and **New** menu options.
3. Select **Volume Contents** as the report type.



The displayed report can be filtered to show one of the following:

- **All Files** - displays all files on the cartridge including deleted files, old versions of files and renamed files.
- **Only Current Files** - displays only the files that can currently be accessed via the Windows file system interface and excludes deleted files, old versions of files and renamed files.
- **Only Deleted Files** - displays only deleted files.

You can further filter the results by specifying a **Search Type (File Name Text Search or Regular Expression Search)**. When **File Name Text Search** is chosen, the search option supports wild cards.

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

Note: A Volume Contents Report will search only on Volumes that have a Volume Contents Catalog file cached on the system. TAR formatted tape cartridges that have not been finalized and have been imported from another XenData6 Server system will not automatically have Volume Catalogs.

10.7.1 Interpreting a Volume Contents Report

An example of a Volume Contents Report is shown below.

Archive Report Generator - [Report2]

File Edit View Window Help

Report Type: Volume Contents
 Volume: 51A765D6-00000000-56028A99
 Showing: All Files
 Search Type: None
 Sorted by: Date

No	File Name	Generation	Version	File Size (bytes)	Date Archived	Type
1	/Test1/01.dat	0	1	1,073,741,824	Sep 23 2015 13:36	Current
2	/Test1/02.dat	0	1	1,073,741,824	Sep 23 2015 13:36	Current
3	/Test1/03.dat	0	1	1,073,741,824	Sep 23 2015 13:36	Current
4	/Test1/04.dat	0	1	1,073,741,824	Sep 23 2015 13:37	Current

Done NUM

The display columns are described below.

- **No** - the sequence number of the file in the display sorted by either date or file name, as defined by the **Sort by** selection.
- **File Name** - 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, the generation number is set to 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 data, XenData6 Server 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 Volume, this column displays the file size stored on the selected cartridge or Volume followed by the total size of the file in bytes.
- **Date Archived** - the date and time the file was archived.
- **Type** The status of the file is displayed as one of the following:
 - **Current** - this is the most recent version of the file, accessible through the archive drive letter.
 - **Renamed** - the file has been renamed and is now accessible under a different name.
 - **Deleted** - the file has been deleted and is no longer accessible except via the **History Explorer**.
 - **Overwritten** - the file has been overwritten and this version is no longer accessible except via the **History Explorer**.

Alert Module

11 Alert Module

Alert Module Overview

The Alert Module is designed for use with XenData6 Server software and is installed separately. It provides e-mail and on-screen alerts that are tailored to the needs of archive system operators, system administrators and IT support personnel. The alerts are derived by filtering and categorizing events recorded in the Windows Event Log.

The Alert Module has two major components:

- **Event Monitor** with integrated e-mail notification that runs on the same server as the XenData6 Server software. The Event Monitor runs an event monitoring service that is pre-configured to detect five different categories of events as they occur in the Event Log. Each Event Category can be mapped to groups of e-mail recipients.
- **On-Screen Messaging** which is a program that may be run on the same server as the Event Monitor or a Windows client. The On-Screen Messaging program can be configured to display via message boxes and system tray notification. The On-Screen Messaging program connects to the event monitoring service on the archive server. If required, the messaging program may be run simultaneously on multiple clients.

The Event Monitor and On-Screen Messaging components are described further in [About the Event Monitor](#) and [About On-Screen Messaging](#), respectively.

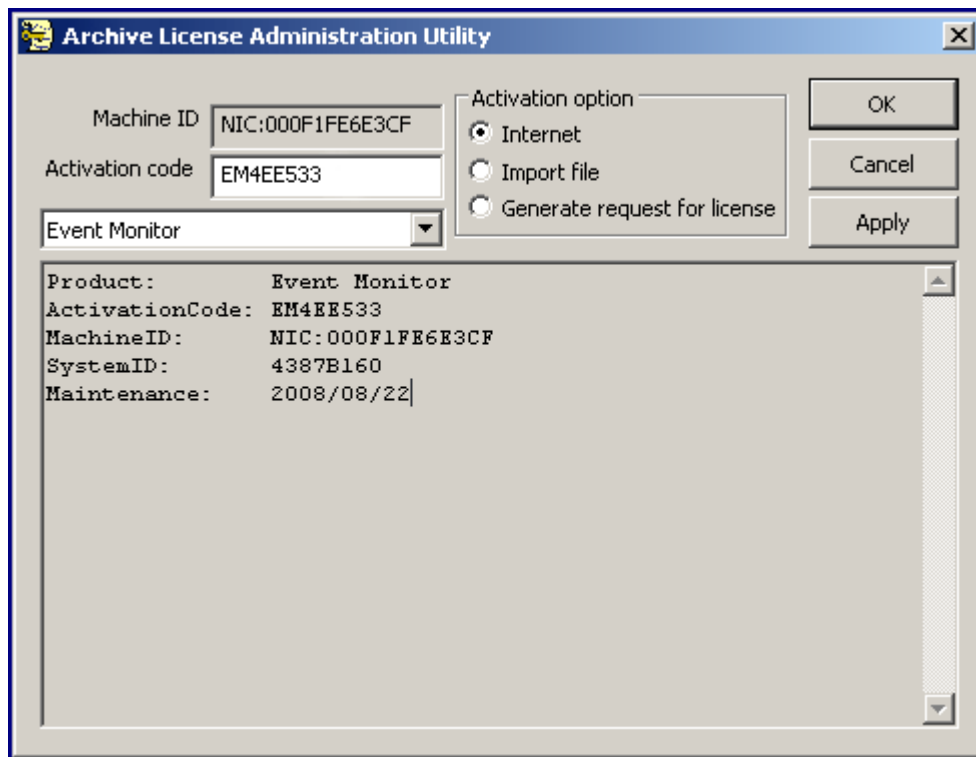
11.1 Installation

Please refer to the separate installation guides for the Event Monitor and On-Screen Notification software modules.

11.2 License Administration

License Administration

The Event Monitor component is licensed using the same License Administration Utility as installed by XenData6 Server software. An Event Monitor activation code is required and is entered into the Activation Code box, as illustrated below.



11.3 Language Support

Language Support

The current version of the Alert Module is available in English.

11.4 Event Monitor About

About the Event Monitor

The Event Monitor runs on the same server as the XenData6 Server software and it provides an event monitoring service with integrated e-mail notification. The event monitor service must be running for correct operation of the On-Screen Messaging program.

The Event Monitor includes a configuration screen that is used to perform the following:

- map categories of events to groups of e-mail recipients, as described in [About Event Categories](#).
- allocate e-mail addresses to groups of e-mail recipients, as described in [About Recipient Groups](#).
- define the e-mail server, e-mail account logon details and e-mail display names, as described in [About the E-mail Server](#).

Set up of the Event Monitor is described in [Event Monitor Configuration](#).

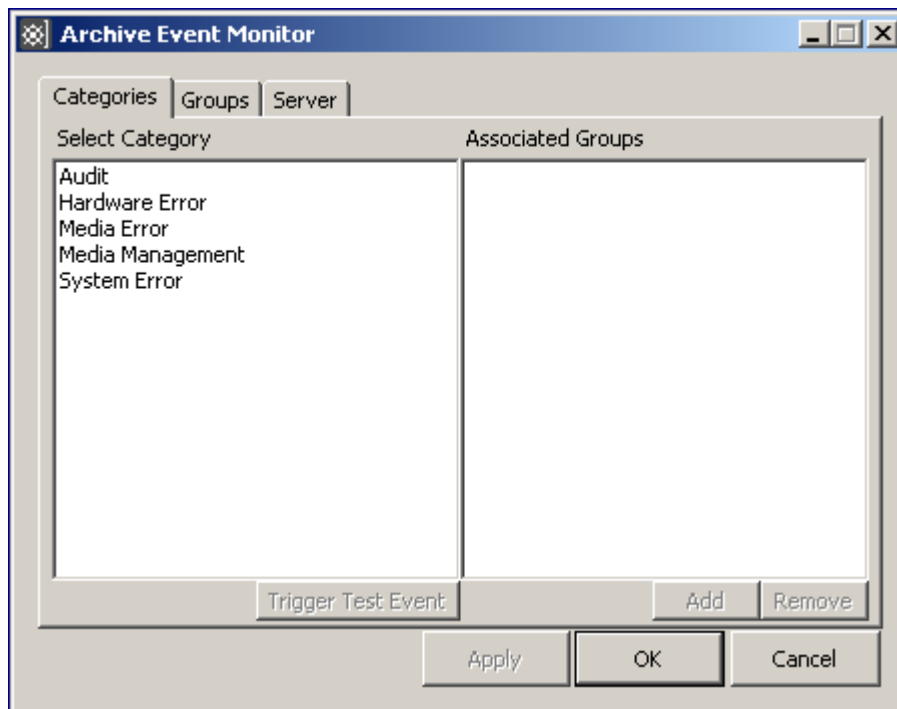
11.5 Event Monitor Configuration

Event Monitor Configuration

The Event Monitor is set up using the configuration program. After initial configuration, changes may be made without need to stop the Event Monitor service.

To start the Configuration Program:

1. Click **Start**
2. Click **Programs**
3. Click **XenData**
4. Click **Configure Event Monitor**



The configuration screen has three tabs as shown above, linked to the following configuration pages:

- Categories. This page is used to map categories of events to groups of e-mail recipients, as described in [Configuring Event Categories](#).
- Groups. This is used to allocate e-mail addresses to groups of e-mail recipients, as described in [Configuring Recipient Groups](#).
- Server. This is used to define the e-mail server, e-mail account logon details and e-mail display names, as described in [Configuring the E-mail Server](#).

After configuration, the event monitoring system can be tested by clicking **Trigger Test Event** on the categories page of the configuration screen. This generates a test event for the selected category. It tests both the e-mail notification and the on-screen messaging (if installed), as it will cause e-mails to be sent to all recipient groups mapped to this category and will initiate an on-screen messages for all connected computers that are running the On-Screen Messaging program.

11.6 About Event Categories

About Event Categories

The Event Monitor is pre-configured with five Event Categories:

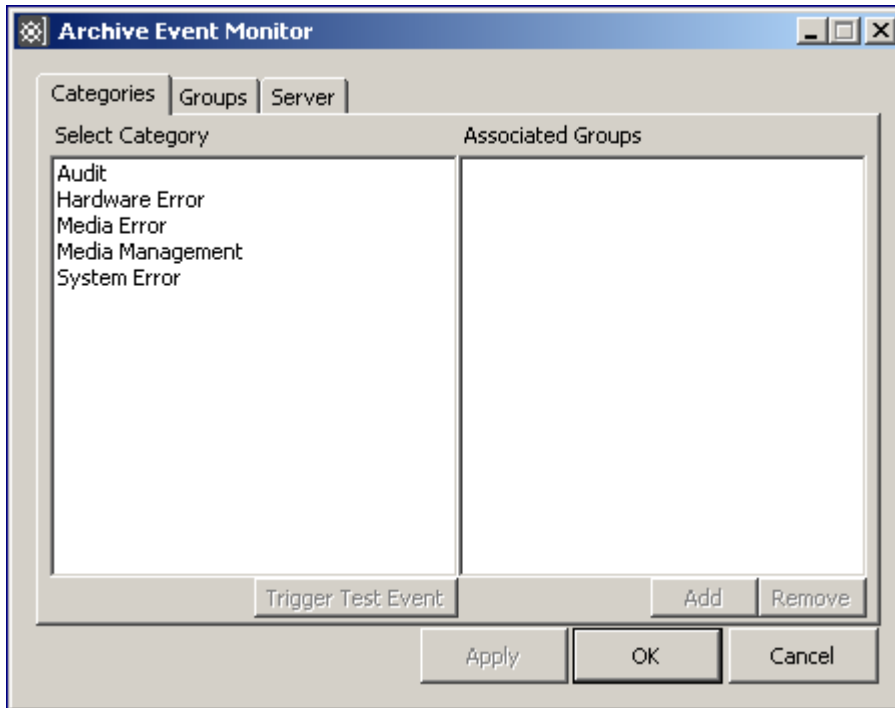
- **Audit** This category of event messages describes the successful completion of media management operations and is particularly useful for the archive operator.
- **Media Management** This category of event messages may require routine action from the archive operator.
- **Media Error** This event category consists of error messages associated with tape cartridge problems.
- **Hardware Error** This event category consists of error messages associated with tape drive and tape library problems.
- **System Error** This event category consists of error messages associated with system problems.

Each Event Category may be mapped to one or more groups of e-mail recipients as described in [Configuring Event Categories](#).

11.7 Configuring Event Categories

Configuring Event Categories

Launch the Event Monitor configuration screen by starting the configuration program as described in [Event Monitor Configuration](#). The configuration screen is shown below.



An event category is mapped to one or more groups of e-mail recipients by using the tabbed Categories page.

To perform mapping of an event category to one or more groups of e-mail recipients:

1. Click on the event category in the left pane
2. Click **Add**, which causes the Add Group display to appear
3. Click to highlight one or more groups in the Add Group display
4. Click **OK**

Repeat this mapping for each event category, as required and then click **Apply**.

11.8 About Groups

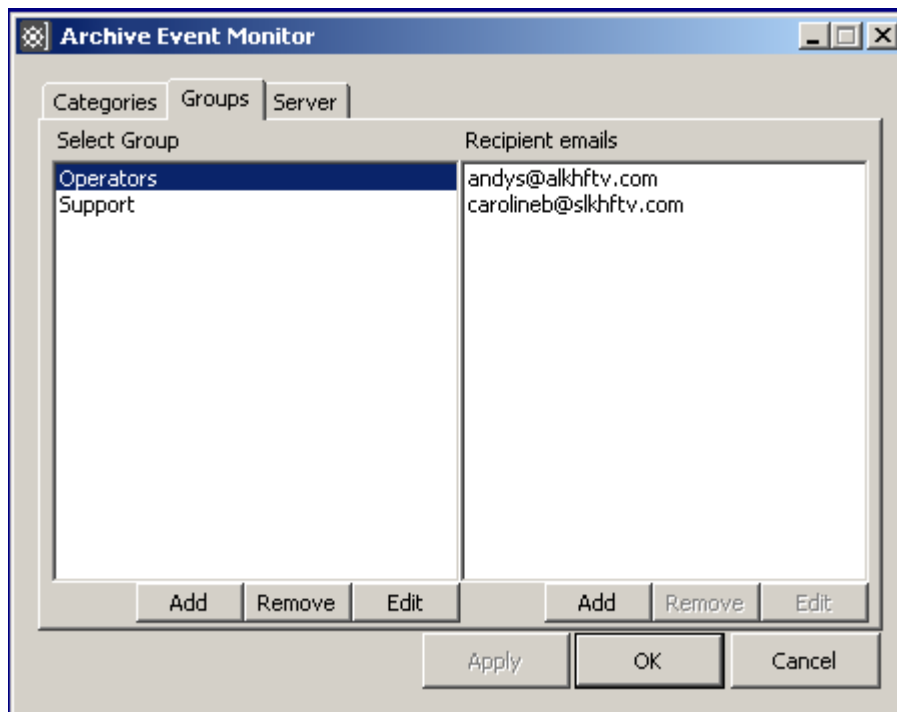
About Recipient Groups

The Event Monitor will send e-mails pertaining to specific event categories to specified groups of e-mail addresses. The groups of e-mail addresses are configured as described in [Configuring Recipient Groups](#).

11.9 Configuring Groups

Configuring Recipient Groups

Launch the Event Monitor configuration screen by starting the configuration program as described in [Event Monitor Configuration](#). Groups of e-mail recipients are configured by using the tabbed Groups page, as shown below.



To add an e-mail address to a recipient group:

1. Click on the group in the left pane
2. Click **Add** under the right pane, which causes the Add email display to appear
3. Enter the e-mail address to be added.
4. Click **OK**

Repeat to add additional e-mail addresses to each group as required and then click **Apply**.

To add a Recipient Group:

1. Click **Add** under the left pane, which causes the Add Group display to appear
2. Enter the name of the group to be added.
3. Click **OK**
4. Click **Apply**

To remove a Recipient Group:

1. Click on the group to be removed in the left pane.
2. Click **Remove**
3. Click **OK**
4. Click **Apply**

To Rename a Recipient Group:

1. Click on the group to be renamed in the left pane.

2. Click **Edit**, which causes the Edit Group display to appear
3. Enter the new name of the group
4. Click **OK**
5. Click **Apply**

11.10 About Email Server

About the E-mail Server

The Event Monitor requires an active e-mail account to send e-mail alerts. The Monitor supports SMTP outgoing servers including Microsoft Exchange Servers and most Internet service provider (ISP) accounts. Popular authentication methods are supported.

Defining the e-mail server and the e-mail account information is described in [Configuring the E-mail Server](#).

11.11 Configuring Email Server

Configuring the E-mail Server

Launch the Event Monitor configuration screen by starting the configuration program as described in [Event Monitor Configuration](#). Defining the e-mail server and configuring the e-mail account is performed by using the tabbed Server page, as shown below.

To define the outgoing (SMTP) server:

In the upper text box enter the DNS address of the SMTP server that will be used to send e-mail and then click **Apply**.

To define the mail account login details:

First, define the authentication method using the drop-down menu options. If further login details are then required (an account name and password), enter them in their respective boxes and then click **Apply**. The authentication types are explained below:

- **None** - No authentication is used when communicating with the server. This requires a server permitting anonymous login, essentially an open relay. (Supported by Microsoft Exchange Server).
- **MD5 Challenge Response** - Authenticate by sending an md5 hash ("fingerprint") of the password when requested by the server, and therefore not requiring the password itself to be transmitted. (Not supported by Microsoft Exchange Server).
- **Basic Authentication (unencrypted password)** - The password is converted into a base 64 number before transmission to the server, but no encryption is used. (This is the most common authentication method which is supported by Microsoft Exchange Server and most ISPs).
- **Plain Text Password** - Both the username and password are transmitted in plain text to the server. This is the least secure method other than no authentication. (Not supported by Microsoft Exchange Server).
- **Windows Authentication** - A Microsoft specific authentication method which uses a user or services logon name and password to authenticate with the server, and therefore no extra authentication is required. (Supported by Microsoft Exchange Server).

To define sender information:

The Sender Name is the display name which will appear in an e-mail client, and the Sender Email is the address which will appear as the 'from address'. Failed-to-deliver e-mail responses will be sent to the 'from address'. Make the required entries in the Sender Name and Sender Email boxes and then

click **Apply**.

To send a test e-mail:

After having defined the outgoing server, mail account login details and sender information, a test email may be sent as follows:

1. Click **Send Test Email**
2. Enter recipient's e-mail address
3. Click **OK**

11.12 Error Reporting

Error Reporting

If the Event Monitor encounters an error associated with sending an e-mail, a message will be added to the Windows Event Log.

Examples of event log messages associated with sending e-mails are given below:

- **No such host is known** - The mail server specified was not found. This means that the mail server address is incorrect.
- **Unexpected ***** response, Last Response: 504 5.7.4 Unrecognized authentication type** - An authentication type is being used which is either unsupported or disabled on the server. Choosing another authentication type may fix the problem. If it does not, it may be necessary to enable the authentication type on the server.
- **Unexpected ***** response, Last Response: 535 5.7.3 Authentication unsuccessful** - Authentication failed, but the authentication type was accepted. This means that the account name/password are incorrect or do not match. Either correct these fields, or set up an account on the server for the desired user. For "Windows Authentication", there must be an account on the server for the account which the service runs under, which may be undesirable, so using another authentication method may be desired.
- **Unexpected RCPT TO response, Last Response: 501 5.5.4 Invalid Address** - The recipient address is invalid. Change the recipient address and try again.
- **Blank sender/recipient address not permitted** - Either the sender or recipient e-mail addresses are blank. Enter an e-mail address for both of these fields to send an e-mail.
- **A connection attempt failed because the connected party did not properly respond after a period of time, or established connection failed because connected host failed to respond** - The connection timed out during communications with the e-mail server. This indicates a problem with the connection to the server or with the server itself. It may be advisable to try another e-mail server until this problem can be resolved.
- **The requested name is valid and was found in the database, but it does not have the correct associated data being resolved for** - An error occurred performing a DNS lookup on the e-mail server address given. It appeared as a DNS entry with no address associated with it. This probably means that the address given is incorrect, although it could mean that the DNS database is out of date (if changes have just been made, and have not propagated yet), or is corrupt (especially if it is a local DNS server).

11.13 On-Screen Messaging About

About On-Screen Messaging

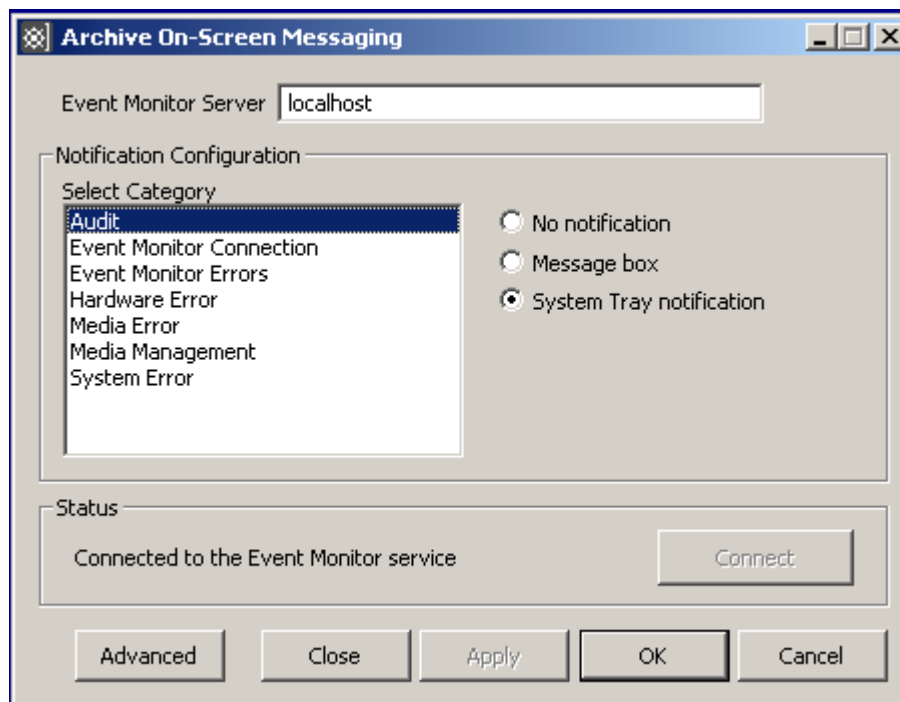
On-Screen Messaging is a program that may be run on the same server as the Event Monitor or a connected Windows client. The On-Screen Messaging program can be configured to display via message boxes and system tray notification. The On-Screen Messaging program connects to the event monitoring service on the archive server and consequently this must be running. If required, the messaging program may be run simultaneously on multiple clients.

Set up of On-Screen Messaging is described in [Configuring On-Screen Messaging](#).


11.14 On-Screen Messaging Configuration

On-Screen Messaging Configuration

The configuration screen for the On-Screen Messaging program is shown below.



For any Event Category, on-screen messaging can be provided via a message box or system tray notification, as shown below.

**To define the Event Monitor Server:**

1. Enter the name of the server running the Event Monitor. (If running on the same server, you may enter 'localhost'.)
2. Click **Apply**

To set up the Notification Configuration for each Event Category:

1. Click on the required Category in the left pane
2. Select either 'No Notification', 'Message box' or 'System Tray notification'

Repeat for each Event Category and then click **Apply**.

To set the notification period for Screen Tray messages:

1. Click **Advanced**
2. Enable the 'Close taskbar notifier automatically' if required.
3. Enter the message retention period in the 'After' box, if applicable.
4. Click **OK**
5. Click **Apply**

After having set up all of the above, connect to the Event Monitor service by clicking **Connect**.

After configuration and connection, on-screen messaging can be tested by clicking **Trigger Test Event** on the categories page of the Event Monitor configuration screen. This generates a test event for the selected category. It tests both the e-mail notification and the on-screen messaging.

Glossary

12 Glossary

Activation Code An Activation Code is required to run the archive management software and enables the chosen library and drive configuration. The [License Administration](#) utility is used to apply an activation code to a system.

Alert Module A companion product which is installed separately. It provides email 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 Windows Event Log.

[Alternate Data Streams](#) are additional named data streams that can be associated with a file. Also called 'Named Streams' and 'NTFS Streams'.

API An acronym for 'Application Program Interface'. Archiving and restoring data with a XenData6 Server system uses normal file system calls issued locally or via a network. Most installations do not use any proprietary XenData6 Server API calls. Software developer kits are available to developers who wish to use proprietary APIs to tightly integrate their applications with XenData6 Server software.

AppleDouble File A term used by Apple to describe how structured files can 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 sometimes termed the 'data fork' and the file with attribute data is termed the 'resource fork'. The resource fork file name is prepended with the characters '._'.

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

Blank Cartridge Set The set of data cartridges in the Archive Management Console which consist of new (unused) cartridges or rewritable cartridges 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 over a network. It is based on the SMB (Server Message Block) network protocol. Apple documentation often refers to the CIFS/SMB network protocol, simply as 'SMB'.

Cache Disk In this documentation, the magnetic or solid state disk volume under control of the archive management software is referred to as the cache disk.

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

[Event Log](#) Part of the Windows Event Viewer within the Computer Archive Management Console. The Event Log is used for XenData6 Server specific messages. If the system is not behaving as expected, the Windows Event Log is the first place that a system administrator should look.

[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 data cartridge spanning.

[File Group](#) A group of files that have the same file management policy and consequently are all treated in the same way by the system (for example, they are all saved to the same Volume Set). Files are assigned to a File Group on the basis of their names.

Finalization A process that involves writing special sequences to a cartridge that prevent further data from being written to the cartridge.

Flushing Files are flushed from a storage device when they are removed to free space on the storage device. XenData6 Server software can be configured to automatically flush files from the cache disk once they are securely stored on tape. After flushing, the file is still visible in the same location in the file system and can still be read so long as the required data cartridges are accessible.

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 A file's generation number is a part of the version control mechanism that XenData6 Server software uses to maintain the complete version history of a file.

History Explorer Enhanced functionality provided within Windows Explorer that is used to obtain the complete version history and status of any file, including deleted and renamed files. It also allows retrieval of all file versions from cache disk or data cartridges.

HSM An acronym for 'Hierarchical Storage Management'.

LTF5 An acronym for 'Linear Tape File System' and one of two [tape cartridge formats](#) supported by XenData6 Server software. 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 and by StorageTek 10000D and later generations.

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 Archive 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 [Archive 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 or solid state disk drives. The XenData6 Server cache disk must be formatted with NTFS.

NTFS Streams See [Alternate Data Streams](#).

Offline See [HSM](#)

Offline File Attribute A file attribute bit (similar to the read-only attribute) defined by Microsoft. XenData6 Server software sets the offline file attribute bit to identify files that have been [flushed](#) from the [cache disk](#). The presence of the offline file attribute is a signal that reading the file may take longer than usual (because the file has to be fetched from the archive). Networking software uses this signal in various ways to ensure that files can be retrieved from storage devices with long access times without suffering network timeouts.

Online See [HSM](#).

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

Primary Replica When [Replication](#) is used, one of the replica cartridges is designated the Primary replica. When new files are written to a replicated Volume, the first copy is written to the Primary replica. A Volume's Primary replica can change; for example if the Primary replica is exported from a robotic library and another replica is available then the other replica will become the Primary. The Primary replica is identified within the Archive Management Console.

Quarantined Cartridge Set A location for cartridges that have been imported into the system 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, unsupported application (such as a backup application) or because the Volume has been [repacked](#).

Repack A XenData6 Server operation which recovers cartridge space used by deleted files and old versions of files by copying "live" data to new cartridges.

SMB See CIFS.

Stand-alone Drive A drive that is not inside a robotic library.

State File An XML file that contains tape information and the Archive Management Console settings, including File Group and Volume Set configuration settings.

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

TAR A term derived from 'Tape Archive' and one of two [tape cartridge formats](#) supported by XenData6 Server software. It is a popular open standard format that defines how file data and file system metadata are written to data tape media. All tape drives support TAR.

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

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

Volume A set of one or more cartridges that have been formatted so that they can contain data. If [Replication](#) is enabled, a Volume consists of a set of replica cartridges which, when up-to-date, all contain identical data.

Volume Contents Catalog Also known as a Volume Catalog. Contains an index of the contents of a specific Volume.

Volume Set A set of one or more Volumes which store 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 Cartridge 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.

Archive Management Console Used to configure all File Group, Volume Set and Volume replication options, and to view diagnostic information about the system.

- A -

Alternate Data Streams 24
Archive Management Console 27
Archive Series Software 7

- B -

Blank Media Set 42

- C -

Cleaning Cartridges 117

- D -

Data Cartridges 53, 53
 Displaying Free Space, Used Space and Overheads 59
 Displaying Information 59
 Exporting 56
 Finalizing 71
 Importing Cartridges 54
 Reformatting 57
 Rewritable Data Cartridges 53
 Setting the Administrator Defined Data 72
 Statistics 61
 Verifying the data on a cartridge 70
 Write Protecting 63
 Writing Files to Data Cartridges 17
Diagnostics 114
Drive Cleaning 116

- E -

Error Management
 Diagnostic Trace Log 114
 System Event Log 113
 Volumes in AlertState 118

- F -

File Groups 29
 Advanced Options 38
 Allocating Files to a File Group 30

 Changing the Order of File Groups 32
 Creating a New File Group 30
 Defining the Path Pattern for a File Group 34
 Editing the Default File Group 29
 Renaming a File Group 31
 Selecting Storage Options for a File Group 35
Folders, Renaming 12

- H -

Hardware 115
 Cleaning Drives 116
 Enabling and Disabling 110

- Q -

Quarantined Cartridge Set 42

- R -

Renaming Folders 12
Rewritable 53

- V -

Verifying data 70
Volume Set 41
 Allocating Media to a Volume Set 47
 Configuring a Volume Set 44
 Creating a New Volume Set 48
 Deleting a Volume Set 52
 Renaming a Volume Set 51

- W -

Windows Explorer Extensions 79
WORM
 WORM Data Cartridges 53
Write Protecting a Data Cartridge 63