

XenData Archive Series Software[®] User Manual

**LTO Edition including Extensions which allow
archiving to ODA devices and object storage**

Version 7.20.3830.0

1	About The Software	8
2	Introduction	10
2.1	Tiered Storage Management	11
2.2	Writing Files to the Archive	12
2.3	Reading Files from the Archive	14
2.4	Functionality Overview	16
2.5	Antivirus Software Compatibility	18
3	LTO and ODA Storage	19
3.1	About Managed LTO Libraries and Drives	20
3.2	About Managed ODA Libraries and Drives	21
3.3	Barcode Management	21
3.4	About LTO Formats	23
3.5	About ODA Formats	23
3.6	About WORM Cartridges	24
3.7	About Rewritable Cartridges	24
3.8	About Cartridge Replication	24
3.9	About LTO Logical Block Protection	25
3.10	About File Fragmentation	25
4	Archiving to Object Storage	26
4.1	About Azure Blob Storage	27
4.2	About S3 Storage	27
4.3	Importing Files Written to Object Storage by Another System	27
4.4	Multi-Site Sync Service	28
4.4.1	Handling of File Versions, Deletions, Overwriting and Renaming	29
4.4.2	How Files are Stored and Synchronized	30
4.4.3	About Application Specific Optimizations	30
4.4.4	Microsoft Office Optimizations	31
4.4.5	Adobe Premiere Pro Optimization	32
5	Global Sync and the S3 Server Interface	34
5.1	Overview of the S3 Server Interface	35
5.2	Synchronizing Multiple Archives	35
6	File Operations, Security and Connectivity	38

6.1	Supported File and Folder Operations	39
6.2	Unsupported Rename Operations	39
6.3	File Version Management	39
6.4	About Offline Files	40
6.5	Handling of Alternate Data Streams	41
6.6	Supported Network Protocols	41
6.7	Free Space Reporting	41
6.8	File Security	42
7	Concepts	43
7.1	About File Groups	44
7.2	About Volumes and Volume Sets	44
7.3	About Volume Catalogs	45
7.4	About Volume Finalization	45
7.5	About Repacking Volumes	46
7.6	About Quarantined Objects	47
7.7	About Pending Write Mode	47
7.8	Partial File Restore and Cartridge Spanning	47
7.9	Offline File Management	48
7.10	Handling of File Delete and Rename Operations	48
8	Administering the System	50
8.1	Tiered Storage Management Console	51
8.2	Configuring LTO and ODA Storage	52
8.2.1	LTO Logical Block Protection	53
8.3	Configuring Azure Storage Accounts	53
8.3.1	Adding Azure Storage Account Access	53
8.3.2	Adding Azure Key Vault Access	55
8.3.3	Configuring a Storage Account	56
8.4	Configuring Amazon S3 Endpoints	57
8.4.1	Adding Amazon S3 Account Access	57
8.4.2	Configuring an Amazon S3 Account	58
8.5	Configuring Wasabi S3 Endpoints	59
8.5.1	Adding Wasabi S3 Account Access	59
8.5.2	Configuring a Wasabi S3 Account	60
8.6	Configuring Generic S3 Endpoints	61
8.6.1	Adding Generic S3 Account Access	61

8.6.2	Configuring a Generic S3 Account	62
8.7	Volume Sets	63
8.7.1	Creating a New Volume Set	64
8.7.2	Renaming a Volume Set	64
8.7.3	Deleting a Volume Set	64
8.7.4	Configuring a Volume Set for LTO or ODA	65
8.7.5	Configuring a Volume Set for Object Storage	66
8.7.6	Configuring Replication for an LTO Volume Set	67
8.7.7	Adding a Volume	68
8.7.8	Displaying Information about a Cartridge	69
8.7.9	Adding User Defined Information for a Cartridge	70
8.7.10	Verifying Cartridges	71
8.7.11	Reformatting Rewritable LTO or ODA Cartridges	72
8.7.12	Exporting Cartridges from an LTO or ODA Library	74
8.7.13	Scanning for Object Storage Containers Created by Other Systems	75
8.7.14	Deleting an Object Storage Container	76
8.7.15	Rebuilding Volume Contents Catalogs	76
8.7.16	Import Folder Structure	77
8.7.17	Import Data	77
8.7.18	Repacking Volumes	78
8.7.19	Cancel Volume Repack	79
8.7.20	Removing Information about a Cartridge from the System	80
8.7.21	Replacing a Missing Replica Cartridge	81
8.7.22	The Blank Cartridge Set	82
8.7.23	Cleaning LTO Drives	82
8.7.24	Displaying Information about Cleaning Cartridges	83
8.7.25	Quarantined Object Set	84
8.7.26	Obtaining Volume Statistics	85
8.7.27	Write Protecting a Volume	85
8.7.28	Finalizing Volumes	86
8.8	File Groups	86
8.8.1	Creating a New File Group	86
8.8.2	Renaming a File Group	87
8.8.3	Changing the Order of File Groups	87
8.8.4	Allocating Files to a File Group	87
8.8.5	Examples of Allocating Files to a File Group	89
8.8.6	Selecting Storage Options for a File Group	90
8.8.7	Selecting a Volume Set for a File Group	91
8.8.8	Selecting File Fragmentation	92
8.8.9	Selecting Disk Retention Rules	92
8.8.10	Changing Disk Retention Rules	94

8.8.11	File Group Advanced Options	94
8.9	Multi-Site Sync Service	96
8.9.1	Adding the Multi-Site Sync Service	96
8.9.2	Adding Cosmos DB Account Access	97
8.10	S3 Server Interface	97
8.10.1	Adding the S3 Server Interface	97
8.10.2	Configuring the S3 Server Interface	98
8.10.2.1	Host	98
8.10.2.2	Access Keys	99
8.10.2.3	Buckets	100
8.10.2.4	Synchronizing Multiple Archives using Shared Volume Sets	101
9	File Explorer Extensions	103
9.1	Flushing of Files and Folders	104
9.2	Pre-fetching of Files and Folders	104
9.3	Smart Copy and Paste	105
9.4	Enhanced Properties	106
9.4.1	Obtaining Enhanced Properties	106
9.4.2	Changing the Object Storage Tier for a File	106
9.5	Volume View	107
9.6	History Explorer	108
10	Metadata Backup	110
10.1	About Metadata Backup	111
10.2	Starting Metadata Backup	111
10.3	Selecting Backup or Restore	111
10.4	Making a Predefined Backup	112
10.5	Making a Custom Backup	114
10.6	Restoring a Backup	118
11	Scheduler	122
11.1	Starting the Scheduler	123
11.2	Adding a Task	123
11.3	The Scheduler Status Display	124
11.4	Editing and Deleting Tasks	125
11.5	Starting and Stopping Tasks	126
11.6	Scheduling Metadata Backup	127
11.7	Scheduling Deferred Write	128

11.8	Scheduling Replication	129
11.9	Scheduling FS Mirror	131
11.10	Scheduling FS Mirror Test	134
12	Reports	138
12.1	Starting the Report Generator	139
12.2	Creating, Saving and Restoring Reports	139
12.3	File Search Report	140
12.3.1	Interpreting a File Search Report	142
12.4	UnArchived Files Report	143
12.4.1	Interpreting an UnArchived Files Report	144
12.5	Volume Contents Report	145
12.5.1	Interpreting a Volume Contents Report	146
12.6	Data Cartridge Contents Report	148
12.6.1	Interpreting a Cartridge Contents Report	149
12.7	Recoverable Space Report	150
12.7.1	Interpreting a Recoverable Space Report	151
13	Alert Module	153
13.1	About the Alert Module	154
13.2	About the Event Monitor	154
13.3	Configuring the Event Monitor	155
13.4	About Event Categories	156
13.5	Configuring Event Categories	156
13.6	About Recipient Groups	157
13.7	Configuring Recipient Groups	157
13.8	About the Email Server	158
13.9	Configuring the Email Server	158
13.10	Error Reporting	160
13.11	About On-Screen Messaging	161
13.12	Configuring On-Screen Messaging	161
14	Diagnostics & Maintenance	164
14.1	Windows Event Log	165
14.2	System Trace Log	165
14.3	Volume Alert State	166
14.4	Library and Drive Diagnostic Information	167

14.5	Cleaning LTO Tape Drives	168
15	System Recovery	169
15.1	Rebuilding a System from Cartridges or Object Storage	170
15.2	In Case of Hardware Failure	172
15.2.1	Options in Case of Library or Drive Failure	172
15.3	Temporarily Disabling LTO or ODA Hardware	172
16	Using Mac Clients	175
16.1	Support of OS X Characters	176
16.2	Hidden File Group Policies	176
16.3	Disabling Alternate Data Streams	177
17	Client Utilities	178
17.1	Installing the Client Utilities	179
17.2	On-Screen Messaging	179
17.3	File Explorer Extensions	179
17.4	Trace File Viewer	179
17.5	FS Mirror Log Files	180
18	Glossary	181

1. About The Software

Archive Series Software

LTO Edition including Extensions which allow archiving to ODA devices and object storage

Version 7.20.3830.0

Copyright © 2001-2021 XenData Limited.

2. Introduction

Version 7 of XenData Archive Series software is available in two editions:

- ❖ LTO Server Edition.
- ❖ Cloud File Gateway Edition.

The LTO Server Edition may be extended by adding software extensions. These include the Optical Disc Archive Extension which adds the ability to archive to Sony Optical Disc Archive drives and libraries, and the Cloud File Gateway Extension which adds the ability to archive to object storage.


This user documentation describes the LTO Server Edition and its optional extensions.

2.1 Tiered Storage Management

XenData Archive Series software runs on a Windows operating system and provides high performance archiving to LTO tape cartridges, Sony Optical Disc Archive (ODA) cartridges, and object storage, including to Azure Blob Storage, Amazon S3 and Wasabi S3. Files are presented in a standard file/folder structure which may be shared over the network. This non-proprietary approach to the interface means file-based applications can write to and read from LTO, ODA or object storage without modification. APIs are also available for interfacing with the archive, but they are not required, avoiding the need to modify every application that uses the archive.

The server running the Archive Series software always includes a dedicated managed disk volume that is used by the XenData software for enhanced performance. The dedicated disk volume is used to store file system metadata, for read and write caching of files and to store files that are to be retained online. It is often termed the 'cache disk', 'disk cache' or 'disk' in this User Manual.

The LTO Server Edition of Archive Series software can be licensed to manage one or more robotic LTO libraries and/or one or more stand-alone LTO drives. After installing the Optical Disc Archive Extension, the software can be licensed to support one or more ODA robotic libraries and/or ODA stand-alone drives. After installing the Cloud File Gateway Extension, with or without the Optical Disc Archive Extension, the Archive Series software can be licensed to support one or more Object Storage accounts. The diagram below illustrates a system with all storage types connected and managed by the software.



The Archive Series software supports three levels of tiered storage:


- ❖ Online with an instance of the file on the disk cache. Online files are read directly from the disk cache. When files are online, there may also be additional instances of the file on LTO, ODA or object storage and, in this case, the timing for how long the file is retained on disk is determined by user-defined disk retention policies.
- ❖ Near-line with at least one instance of a file on an available LTO cartridge, ODA cartridge or Object Storage account and no instance on the cache disk. The LTO or ODA cartridge will be in an attached robotic library or a stand-alone drive. When near-line files are read, they are restored from the LTO, ODA or object storage. Disk retention policies determine how long a file will be retained on disk cache after it has been read.
- ❖ Offline with no instance on the cache disk and one or more instances on LTO or ODA 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 user of the identities of the cartridges that contain the data.

Regardless of the number of LTO or ODA robotic libraries, the number of LTO or ODA stand-alone drives and the number of Object Storage accounts managed by the Archive Series software, all files in the archive appear within one Windows logical drive letter on the server whether those files are on the disk cache, on near-line or offline storage.


2.2 Writing Files to the Archive

The file system managed by Archive Series software appears as a single logical drive letter. This may be accessed as a network share or by an application running on the same computer as the

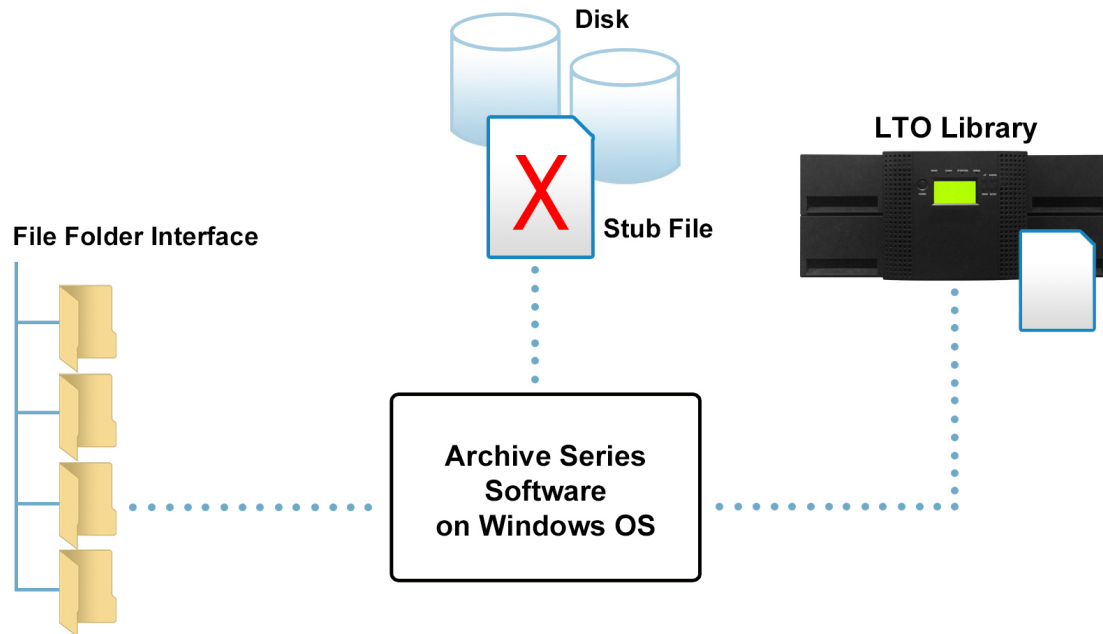
Archive Series software. When files are written to the system, they are always first written to the cache disk as illustrated here for a system with a single LTO library.



Once the file copy to disk is complete it will then be written to Archival Storage (LTO/ODA cartridges or Object Storage).




When a file has been written to its designated locations, it becomes eligible for flushing from cache. After flushing, the full file is no longer retained on the disk cache. The flushed file has all the same properties as the original except the Microsoft offline attribute is set indicating that the full file is no longer immediately available. The flushing operation frees up space on the cache disk because the file representation on the disk is just a few Kilobytes in size.



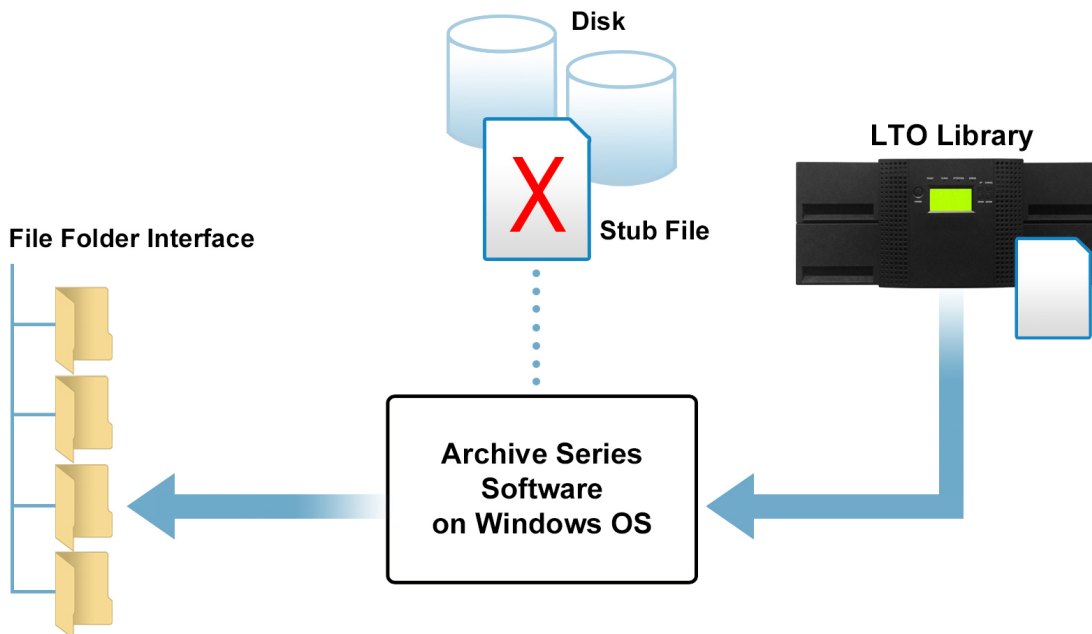
Flushing the file may be scheduled to occur immediately after the file has been successfully written to LTO, ODA or object storage. Alternatively, it may be scheduled to occur a defined time after the file was written or last read. Another option is that some files may be retained on disk cache permanently, as well as being written to Archival Storage. The rules that determine how long the full file will be retained on disk are defined by using the Tiered Storage Management Console. The retention rules may be different for different file types and for different folders.

2.3 Reading Files from the Archive

When it comes to reading files, they are restored seamlessly whether the file is held on the disk cache or near-line LTO, ODA or on Object Storage. In the case shown here, where there are instances of the file on both the disk and LTO, ODA or Object Storage, the file is simply restored from disk.



When the file has been flushed from the cache disk and the full file is on near-line LTO, ODA or object storage, it is restored directly and automatically from the Archival Storage.



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 Archive Series software puts a message into the Windows Event Log that identifies which data cartridges contain the file.

When the XenData Alert Module is configured, on-screen messages and e-mail alerts are also generated that identify the file name and the cartridges that contain the file.

2.4 Functionality Overview

- ❖ **Standard File Interface** The system running Archive Series software accepts all file types – from an MXF to a WORD document - and presents them in a single Windows file/folder structure. Files are written to and retrieved from the archive as though from a standard disk-based volume or network share.
- ❖ **Standard Network Protocols** In addition to supporting CIFS/SMB, NFS and FTP network protocols, it may be accessed locally as a logical drive letter.
- ❖ **Remote Access** Adding an S3 Server Interface allows the archive to be accessed as object storage. Using this interface, files may be written to and read from the archive from anywhere worldwide using fast and secure HTTPS. This functionality requires an additional license subscription; it is described in [Overview of the S3 Server Interface](#).
- ❖ **Manages a Disk Cache**, The user defines policies for disk caching that can be tailored for different file types and folders, allowing frequently accessed files to be retained on disk.
- ❖ **Supports LTFS and TAR tape formats** This avoids proprietary formats and vendor lock-in.
- ❖ **Supports WORM LTO and ODA Cartridges** This is ideal for compliance applications.
- ❖ **Self-Describing LTO and ODA Cartridges** Each LTO or ODA cartridge contains all the file system metadata necessary to recover all the files stored on it. This allows individual LTO and ODA cartridges to be easily transferred between archive systems.
- ❖ **LTO Cartridge Replication** The software automatically generates replica LTO cartridges that may be exported from the library for off-site retention.
- ❖ **Dynamic Expansion of LTO and ODA Cartridge Groups** The system will dynamically expand LTO and ODA cartridge groups to meet capacity demands. This means that the system runs automatically with an LTO or ODA library without need for human intervention.
- ❖ **Checksum Verification** Archive Series software implements an automated checksum operation for all data written to LTO, Azure Blob Storage, Amazon S3 and Wasabi S3 .
- ❖ **LTO Cartridge Spanning** User defined policies can be set to allow or prevent files being spanned across multiple LTO cartridges. Additionally, the transfers of multiple files and folders will be automatically spanned across multiple cartridges.

- ❖ **Manual Pre-Fetch and Flush of Files from Disk Cache** The pre-fetch operation creates an instance of a file on the disk cache copying it from LTO, ODA or Object Storage. The flush operation removes a file's data from the disk to free space, replacing it with a stub file.
- ❖ **Optimized Restores** The system restores a queue of files in the shortest possible time from LTO and ODA. The restore requests are processed in an order that minimizes unnecessary tape movement and ODA disc swapping. This greatly decreases total restore time when restoring multiple small files from LTO or ODA.
- ❖ **Partial File Restores** With very large files there is often a need to read only a portion of the file. For example, this frequently occurs with multi-gigabyte video files when a short clip is requested. The Archive Series software supports partial file restore (PFR) from LTO, ODA and Object Storage based on file restore requests with specified byte ranges. A companion product, XenData Workflow API, extends this functionality to provide PFR based on timecode ranges.
- ❖ **Migration from one Storage Type to Another** A repack function allows files stored on one storage type to be migrated to another without changing the file name and path. Examples are migration from one generation of LTO to a later generation; migration from LTO to Object Storage.
- ❖ **Recovery of space from LTO and ODA** Rewritable LTO and ODA cartridges may be repacked to recover space from deleted and overwritten files.
- ❖ **File Version Control** The software provides comprehensive file version control when files are stored on LTO or ODA. This means that deleted files and old file versions may be restored from LTO or ODA (unless the files have been purged using a repack operation).
- ❖ **Local File-Folder Sync** FS Mirror utility synchronizes any file-folder structure on an accessible network share or on the local computer to the archive volume and vice versa. This is functionality that is licensed separately and is described in [Scheduling FS Mirror](#).
- ❖ **Global Sync of Multiple LTO or ODA Archives** Adding an S3 Server Interface and a XenData synchronization service links multiple LTO archives, creating a single global file system accessible anywhere worldwide. This functionality requires additional license subscriptions; it is described in [Global Sync and the S3 Server Interface](#).
- ❖ **Metadata Backup and Restore** A file system metadata backup and restore utility is provided which provides rapid system restore in case of rebuild after disk cache failure.

- ❖ **Alert Module** A software module is included which provides e-mail and on-screen alerts. These are tailored to the needs of system operators, system administrators and IT support personnel.
- ❖ **System Reports** A Report Generator allows you to create, save and restore a range of different reports about the files managed by the system.
- ❖ **Industry Standard File Security** The tiered storage managed by Archive Series software integrates fully with the Microsoft Windows security model based on Active Directory. The solution can be installed within a domain or workgroup.

When using the Cloud File Gateway Extension, the following functionality is available.

- ❖ **Supports Import of Objects written by 3rd Party Applications** Containers or buckets created by 3rd party applications such as Azure Storage Explorer, AzCopy or Wasabi Explorer may be imported into the gateway file system.
- ❖ **Scheduler Optimizes Internet bandwidth** Time windows are scheduled so applications can write to the disk cache while postponing a copy being made to Object Storage, allowing Internet bandwidth to be optimized when in high demand.
- ❖ **Encryption** All data transferred between the Object Storage and the Cloud File Gateway installation employs the HTTPS communication protocol, using secure socket layer (SSL) encryption.
- ❖ **Multi-Site Sync** Multiple Cloud File Gateway instances may be synchronized using XenData's Multi-Site Sync service. It provides a global file system accessible from each gateway and the networks attached to them. This is ideal for worldwide collaboration and file sharing across multiple sites. The Multi-Site Sync service provides immediate synchronization of the global file system across all gateways.

2.5 Antivirus Software Compatibility

When installing anti-virus protection on the computer running the Cloud File Gateway software, it is important to choose an anti-virus (AV) solution that has been certified. The XenData 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 XTN1801 available in the support section of the XenData [website](#).

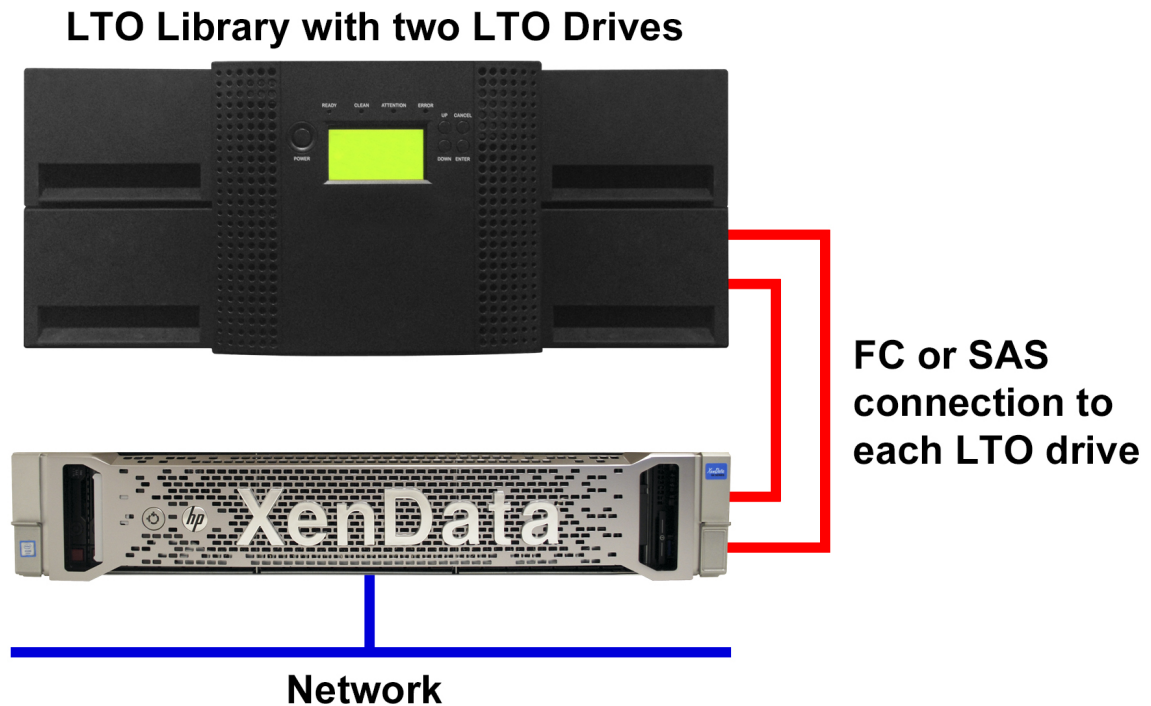
Please check this guide to ensure that you are installing the XenData software on a machine that meets the installation prerequisites.

3. LTO and ODA Storage

LTO (Linear Tape Open) tape cartridges provide high capacity data storage with high data transfer rates. The cartridges are very stable with a shelf life of 30 years. ODA (Optical Disc Archive) cartridges from Sony are an alternative on-premises storage option with faster access times than LTO and a 100 year shelf life.

3.1 About Managed LTO Libraries and Drives

XenData Archive Series software may be configured and licensed to manage one or more robotic LTO libraries, optionally combined with one or more stand-alone LTO drives. The software supports a wide range of LTO library models. A Windows server running Archive Series software with a tape library is illustrated in the diagram below.



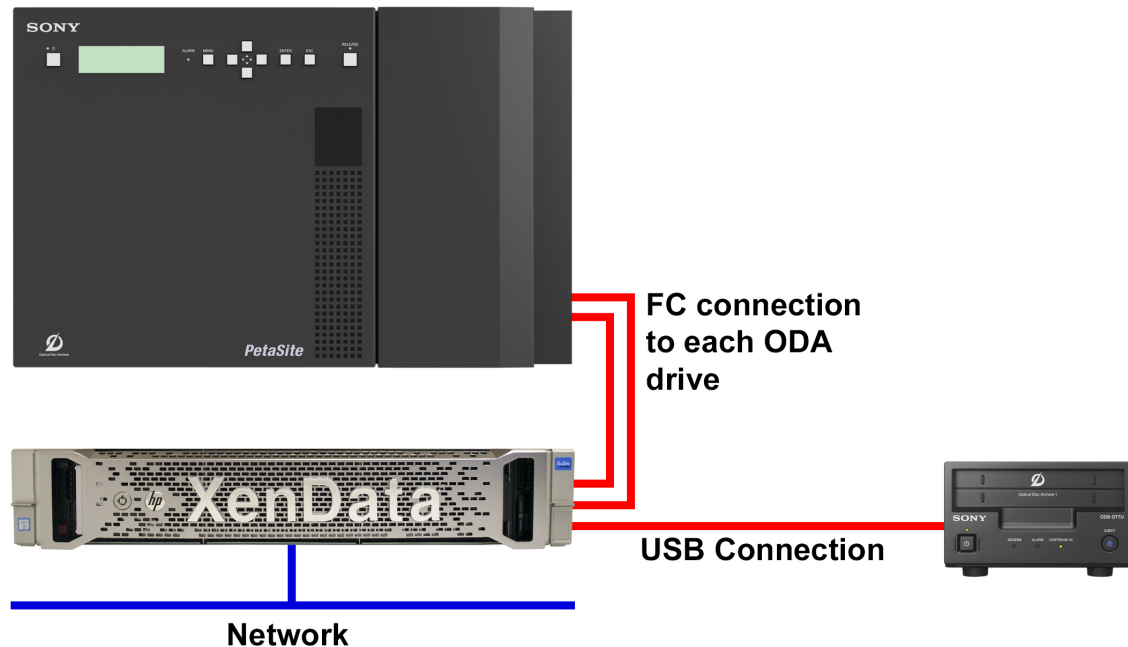
LTO libraries are typically connected via Fibre Channel (FC) or Serial Attached SCSI (SAS) to the server running Archive Series software. FC libraries may be connected directly to FC ports on the archive server or via a fibre channel switch.

LTO drives, whether internal to a library or stand-alone units, need [cleaning using a cleaning cartridge](#) from time to time. For LTO libraries, the Archive Series software will automatically perform drive cleaning as needed.

3.2 About Managed ODA Libraries and Drives

XenData Archive Series software may be configured and licensed to manage one or more Sony ODA libraries, optionally combined with one or more stand-alone ODA drives. A Windows server running Archive Series software with a Sony ODA library and stand-alone ODA drive is illustrated in the diagram below.

Sony ODA Library with two ODA Drives



The Sony FC libraries may be connected directly to FC ports on the archive server or via a fibre channel switch. Sony ODA stand-alone drives are connected via USB.

Unlike LTO drives, Sony ODA drives do not need cleaning.

3.3 Barcode Management

Barcode labels are available for all LTO and ODA data cartridge formats supported by XenData Archive Series 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.

All LTO and ODA cartridge types include an in-cartridge memory chip. When a cartridge is used in a robotic library, Archive Series 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 Archive Series 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 Tiered Storage Management Console, Event Log, History Explorer and Report Generator. In addition to these functions, the Archive Series software matches barcodes for replicated LTO cartridges and selects the cartridges 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 of the barcode character positions. 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 have the following sequence of barcodes in the Blank Cartridge Set

- ❖ X0007AL8
- ❖ X0008AL8
- ❖ X0008BL8
- ❖ X0009AL8
- ❖ X0009BL8

the system would next allocate X0008AL8 and X0008BL8 to a replicated set of tapes. However, if the system were allocating a cartridge to a non-replicated set, it would select X0007AL8.

LTO libraries use a barcode format called '3 of 9' or code 39. The last two digits of the human readable format represent the data cartridge format. For example, L8 is used for LTO-8 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 LTO cleaning cartridges have a specific barcode label format which always starts with CLN. This allows Archive Series software to recognize a cleaning cartridge without putting it in a drive and unnecessarily using a cleaning cycle.

3.4 About LTO Formats

XenData Archive Series software can format [rewritable](#) LTO-5 cartridges and later generations in LTFS. All generations later than LTO-2, whether [WORM](#) or rewritable, can be formatted with TAR.

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. XenData Archive Series software extends the TAR format by maintaining a catalog (the Volume Catalog) that includes an index of the cartridge contents. When the tape is full, the Volume 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 XenData Archive Series 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. 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.

3.5 About ODA Formats

Sony Optical Disc Archive cartridges use only one format, no matter what software is initializing and writing to the cartridge. This format is based on the UDF optical disc file system format and is used by XenData Archive Series software.

3.6 About WORM Cartridges

WORM (Write-Once, Read Many) storage is ideal for legal compliance applications. It is also suitable for long term storage of data that will not change.

A WORM LTO cartridge is identical to a rewritable LTO cartridge of the same generation with the following exceptions: the cartridge memory identifies it to the drive as WORM, the servo tracks are slightly different to allow verification that data has not been modified and the bottom half of the cartridge shell is gray. LTO drives recognize WORM cartridges and prevent reformat. This ensures that the data written to a WORM LTO cartridge cannot be erased and re-written.

A WORM Optical Disc Archive (ODA) cartridge uses a different recording medium to a rewritable ODA cartridge. The recording layer is fundamentally write-once in nature. This intrinsic characteristic of WORM ODA cartridges ensures that written data cannot be erased and re-written.

XenData Archive Series software allows files to be overwritten and deleted even when using WORM cartridges. The deleted and overwritten files continue to be accessible from the cartridges using [History Explorer](#). However, only files that have not been deleted and the latest versions of files are available from the Archive Series file-folder interface.

3.7 About Rewritable Cartridges

A rewritable LTO and ODA cartridge may be reformatted which erases all data written to the cartridge and prepares it for reuse.

3.8 About Cartridge Replication

XenData Archive Series software can automatically create multiple LTO cartridge replicas. Automatic cartridge replication is not available for ODA cartridges.

By having two or more copies of every LTO cartridge, replication provides additional data protection. Best practice when duplicating cartridges is as follows:

- ❖ Enable automatic replication using the Tiered Storage Management Console
- ❖ After a replica set of LTO cartridges becomes full, export either of the LTO cartridges from the library
- ❖ Retain the exported replica LTO cartridges in an offsite location

Note that 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.

3.9 About LTO Logical Block Protection

End to End Logical Block Protection is functionality introduced with LTO-5 drives that allows XenData Archive Series software to cooperate with the drive to ensure the highest level of data integrity. The XenData software calculates a checksum for each block of data as it reads a file and sends the checksums to the LTO tape drive alongside the data. The LTO drive performs a read-after-write verification operation and calculates its own checksum from the data stored on the tape, then compares the two checksums to ensure that the data on the tape matches the original data.

With LTO-5 and LTO-6 drives, there is a significant CPU overhead to perform the calculation which may reduce writing speed. From LTO-7 onwards, an alternative form of CRC was added which is used by Archive Series software and introduces no significant increase in CPU utilization.

3.10 About File Fragmentation

XenData Archive Series software optionally splits a file written to LTO into chunks. This optional file fragmentation should be enabled to support partial file restore (PFR) from LTO. Furthermore, file fragmentation allows extremely large files to be spanned across multiple LTO cartridges.

If file fragmentation is enabled, the system has the following characteristics:

- ❖ When a portion of a file is read from LTO, 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 LTO 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 LTO.
- ❖ On writing a file, the data may span multiple LTO 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 for an LTO Volume Set, the system has the following characteristics:

- ❖ When a file or portion of a file is read from LTO, the whole file will be restored
- ❖ When a file is modified, the new version of the file will be completely written to LTO
- ❖ On writing a file, the system will always write the whole file to a single LTO Volume

Note that partial file restore is supported with ODA and Object Storage Volume Sets which are not fragmented.

4. Archiving to Object Storage

Adding the Cloud File Gateway Extension allows the XenData Archive Series software to archive to object storage as an alternative to LTO or ODA. Supported object storage includes Azure Blob storage, Amazon S3 and Wasabi S3. The Cloud File Gateway must be licensed to support the required Object Storage capacity.

The Cloud File Gateway Extension is complimentary to the S3 Server Interface described in [Overview of the S3 Server Interface](#), which allows the entire archive to be accessed as object storage.

4.1 **About Azure Blob Storage**

The Microsoft term for a single Object Storage asset is a Blob. Each file stored in Azure Blob Storage by Archive Series software is written to one Blob. Blobs are grouped in Containers within an Azure Blob Storage account. The XenData Archive Series software uses the term 'Volume' to refer to an Azure Container. A Volume Set is a set of one or more Volumes which store files from designated File Groups.

When 1 million Blobs have been written to a Volume, i.e. to a Blob Container, it is identified as full. The system will create a new Volume automatically. The creation of the new Volume and its use for new data are completely automatic. Consequently, the cloud Object Storage will continue to expand automatically as more capacity is required.

4.2 **About S3 Storage**

Simple Storage Service (S3) is an Object Storage service, originally developed by Amazon Web Services. Each file stored in S3 by Archive Series software is written to one Object. Objects are grouped in buckets within an S3 storage account.

The XenData Archive Series software uses the term 'Volume' to refer to an S3 buckets. A Volume Set is a set of one or more Volumes which store files from designated File Groups.

When 1 million Objects have been written to a Volume, i.e. to an S3 Bucket, it is identified as full. The system will create a new Volume automatically. The creation of the new Volume and its use for new data are completely automatic. Consequently, the Object Storage will continue to expand automatically as more capacity is required.

4.3 **Importing Files Written to Object Storage by Another System**

Files written to accessible Object Storage containers or buckets by compatible non-XenData applications may be imported into the gateway file system where they will be read-only

accessible. This is performed using a scriptable sequence of operations described in the next paragraph. Examples of compatible non-XenData applications are Azure AzCopy, Azure Storage Explorer, AWS CLI and Wasabi Explorer.

After files have been uploaded to a new container or bucket by another system, they appear in the file-folder interface and become read-only accessible by performing the following steps:

- ❖ Scan the Object Storage account for new Volumes, as described in [Scanning for New Volumes](#).
- ❖ Build a contents catalog for the new Volume, as described in [Rebuilding a Volume Contents Catalog](#)
- ❖ Use Import Folder Structure or Import Data, as described in [Importing Folder Structure](#) and [Importing Data](#)


A scheduled PowerShell script may be used to perform these steps automatically.

If the Object Storage container is an Azure Blob, you can forgo the above steps thanks to the Azure Change Feed. With the Azure Change Feed, file creations, updates and deletes on a connected Azure Blob Storage container are quickly replicated on your XenData system, without the need for manual intervention. The Azure Change Feed requires a separate XenData service, which can be found in the Cloud File Gateway, or Cloud File Extension installers.

4.4 Multi-Site Sync Service

XenData's Multi-Site Sync service may be used with two or more Cloud File Gateway instances that share access to cloud object storage. The service shares the file-folder structure created by each gateway with all gateways.

When one gateway writes a file to cloud object storage, it immediately appears as a stub file on all other gateways. The example below illustrates a Multi-Site Sync system with three gateways: one in London, one in New York and another in Los Angeles.



As soon as a file is written to the cloud by the New York Gateway, it immediately appears as a stub file in the other locations and is available for restore.

Multi-Site Sync implements global file locking for file types that have application optimizations, as described in later sections of this guide.

4.4.1 Handling of File Versions, Deletions, Overwriting and Renaming

The Multi-Site Sync service controls which files will be shown in the globally shared file system. It always presents the latest version of a file whichever gateway wrote it.

File deletion, overwriting and renaming are handled as described below.

- ❖ **Overwriting Files** A file can be overwritten in the global file system by any gateway, even if it was not written by that gateway. When a file is overwritten by a gateway that did not originally write it, the original version is retained permanently in the cloud object storage and may be restored using the XenData History Explorer function.
- ❖ **Deleting Files** A file may only be deleted by the gateway that last wrote it.

- ❖ **Renaming Files** When the Multi-Site Sync service has been installed on a gateway, renaming of files is no longer supported, except when permitted by an application specific optimization.

4.4.2 How Files are Stored and Synchronized

Each gateway has its own set of containers or buckets, called a Volume Set, to which it has read-write access. Additionally, each gateway has read-only access to all other Volume Sets i.e. to containers or buckets written by the other synchronized gateways.

The XenData Multi-Site Sync service uses Azure Cosmos DB, Microsoft's globally distributed, low latency database service to perform the synchronization. The service is not limited to Azure; it is applicable to all Object Storage providers supported by the Cloud File Gateway.

When a file is written to the system, it will be written as a single object to the Volume Set owned by the gateway that wrote it. As soon as the file has been fully written to object storage, the Sync service will immediately update all other gateways and a stub file will be shown in each of their file systems. The file may then be restored by any gateway simply by reading the stub file.

When a file is overwritten by a gateway that did not write it, the file will be written to the Volume Set owned by the gateway performing the overwrite. Metadata for the overwritten version of the file is then immediately propagated throughout in the global file system and the new version is available to all gateways. If a gateway had previously cached the prior version of the file on its managed local disk, that prior version will be immediately flushed and the file system will show a stub file of the latest file version.

4.4.3 About Application Specific Optimizations

During the installation of Cloud File Gateway software, application specific optimizations may be selected. These optimizations enhance compatibility with applications that use complex sequences of operations to save files to a gateway and allow users to work collaboratively over the global file system. They address two aspects of file interaction: simplification of complex save operations and management of owner files.

- ❖ **Simplification of Complex Save Operations**

Some applications use complex save sequences that create and rename multiple intermediate files. Complex saves are performed by applications for reliability reasons; they ensure that writing a new version of a file is resilient to any glitches in the storage process. However, if allowed to propagate across the entire global file system, they would result in a lot of unnecessary file changes including operations (such as renames) that are not typically

supported by cloud storage providers, which in turn could result in unreliable synchronization.

The XenData application-specific optimizations detect the intermediate files and file renames created by an application and save these locally. It only synchronizes the final file system changes throughout the global file system. Consequently, it maintains the reliability associated with an application's complex save process but simplifies the changes that are propagated throughout the global file system, further improving reliability.

❖ **Management of Owner Files**

Some applications use 'owner files' to support collaborative work over a network. An owner file is a small temporary file that identifies that the file has been opened by another user and includes the name of that user. These are sometimes called 'lock files' but we will use the term 'owner files' to avoid confusion with file locking which refers to a specific file state.

The XenData application optimizations detect the creation of an owner file and store the user data within the Cosmos DB database. Owner files are rapidly created and synchronized throughout the global file system. This means that applications running across all gateways and the networks attached to the gateways are presented with owner files as though all users were on one local area network.


4.4.4 **Microsoft Office Optimizations**

The Microsoft Office optimizations apply to the following file types:

- ❖ docx
- ❖ xlsx
- ❖ pptx

They are limited to current versions of Word, Excel and PowerPoint running on Windows client computers that access the gateway via a local network.

An example of a message displayed by Word when opening a file that has already been opened by another user is shown below.




4.4.5 Adobe Premiere Pro Optimization

Adobe Premiere Pro project files have an extension 'prproj' and when a project file is opened by a user, a temporary owner files is created that has a 'prlock' extension. The XenData Premiere Pro optimization applies to these files.

Premiere Pro can create a team project which does not have a project file; the project data lives in Adobe's Creative Cloud. However, there is usually a need to convert between team projects and standard projects, at least from time to time. Even if primarily using team projects, the Adobe Premiere Pro optimization should be installed.

The Adobe Premiere Pro optimization is certified for use with Premiere Pro 2021 running on Mac and Windows machines.

When a user opens a project file that has already been opened by another user on the same or a different gateway, the Premiere Pro user interface will identify it with a lock icon and will display the user that has it open, as shown below.

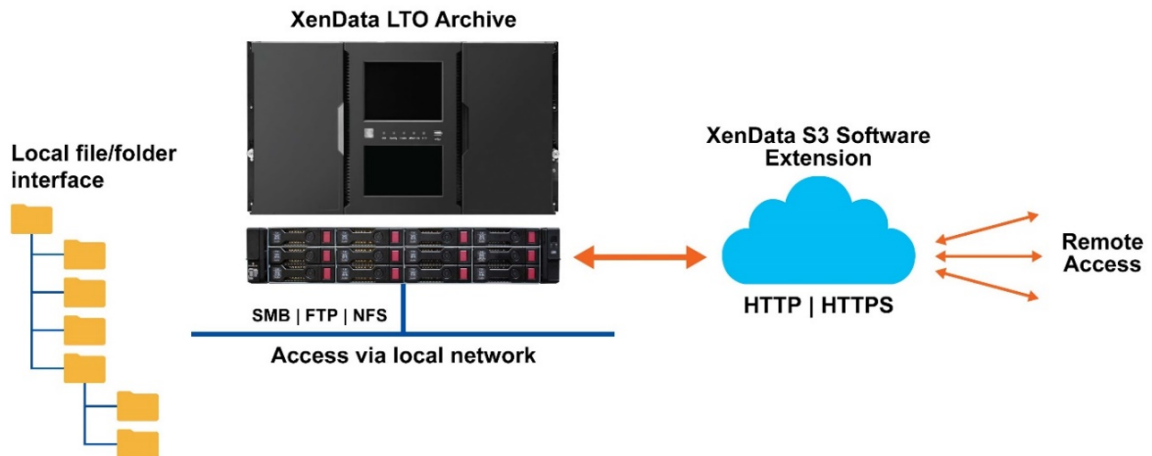


As for all projects with an Adobe lock, it will not be possible to save the project file as this options will be greyed but 'save as' will be available.

5. Global Sync and the S3 Server Interface

5.1 Overview of the S3 Server Interface

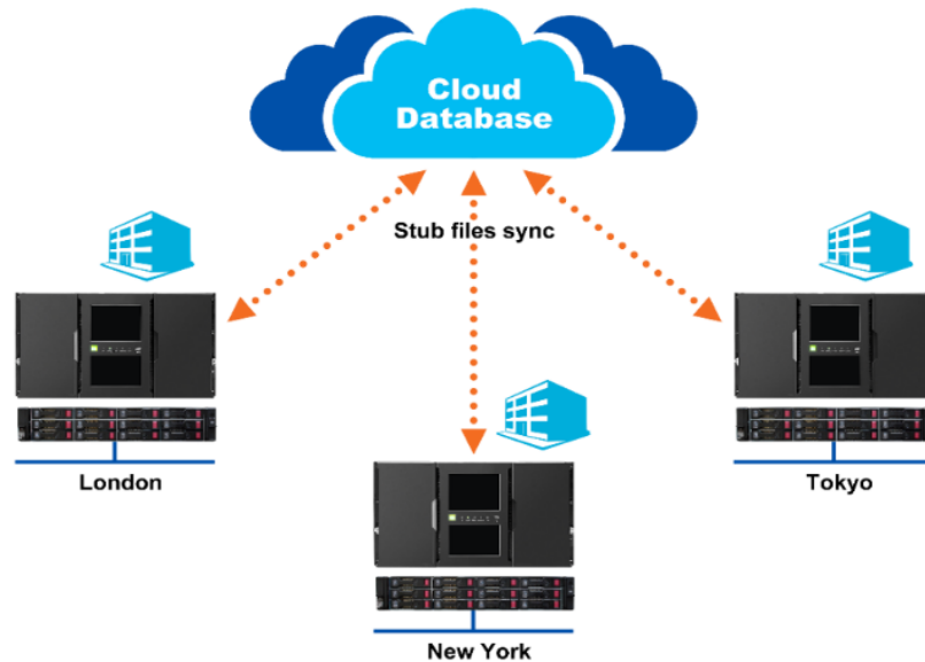
The S3 Server Interface provides ‘front end’ access to the archive, allowing files to be written to and read from the archive from anywhere worldwide using fast and secure HTTPS. The on-premises interface is not affected: the archive continues to be accessible locally as a logical drive letter and over the local network using SMB, FTP or NFS. This is illustrated below for an archive that is configured with an LTO library.



The S3 Server Interface is complementary to the XenData Cloud File Gateway software, described in [Archiving to Object Storage](#), which provides ‘back end’ access to object storage.

5.2 Synchronizing Multiple Archives


Global Sync is a XenData synchronization service that links multiple archives, creating a single global file system accessible anywhere worldwide. Each archive system must have the S3 Server Interface enabled. When a user makes a change by writing, overwriting or deleting a file in one location, that change is propagated to all locations.



The XenData synchronization service uses a low latency Cosmos DB database running in Azure. As soon as a file is archived to LTO, the cloud database is updated and the file becomes available as a stub file within the global file system.

Global Sync requires minimal Internet bandwidth for synchronization across all locations because only stub file representations are synced. Even though the Global Sync service uses a cloud database, the files themselves are never stored in public cloud object storage, avoiding cloud storage and egress fees.

When a stub file is read and when the full file is stored at another location, the file is transferred directly using peer to peer multi-threaded HTTPS. This delivers secure fast file transfers.



The administrator at each location defines which Volume Sets will be included in the globally synchronized file system.

Each archive system includes a disk cache that is managed by the XenData software. Consistency with the global file system is ensured even if files from other LTO archives have been cached locally. The cached files are changed to ensure consistency with the global file system. For example, when a file originally archived in London has been cached in Tokyo and then the file is overwritten by the London archive, the file cached in Tokyo will be replaced with a stub file that links to the latest version.

Global Sync allows an organization with multiple locations to link its LTO archives into one synchronized system. And it provides a consistent up-to-date set of files across the entire distributed organization.

6. File Operations, Security and Connectivity

The XenData Archive Series software is tightly integrated with the Windows operating system and supports most file and folder operations. It is fully compliant with the Microsoft security model.

6.1 Supported File and Folder Operations

You can write, read, delete, overwrite and rename files. You can create new folders, rename empty folders and delete empty folders. The system supports partial file restores which means that when an application sends a request to read only a specific byte range from within a file, only that portion of the file and not the whole file is restored.

Note for Users familiar with the LTO Server Edition: In the case of restores from LTO, file fragmentation must be enabled, and only file fragments that contain the requested byte range will be read. File fragmentation is not used for Object Storage and is not required to enable partial file restore.

6.2 Unsupported Rename Operations

The Archive Series software does not support renaming folders after a file has been added to that folder.

When using the optional S3 Server Interface, renaming of files is also not supported.

6.3 File Version Management

The standard 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. Archive Series software maintains a complete version history of files written to LTO or ODA. 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 the software. The default configuration of Archive Series software does not maintain version history when writing to Object Storage; in this case, when files are over-written or deleted, the old versions or deleted files are removed from the Object Storage.

The file version numbering convention is as follows.

- ❖ When a file is initially created it has a version number of 0. Version 0 does not contain any data; it 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 this option has been configured in the Tiered Storage Management Console. If the file is subsequently re-opened and has more data written to it, the version number will once again be incremented.

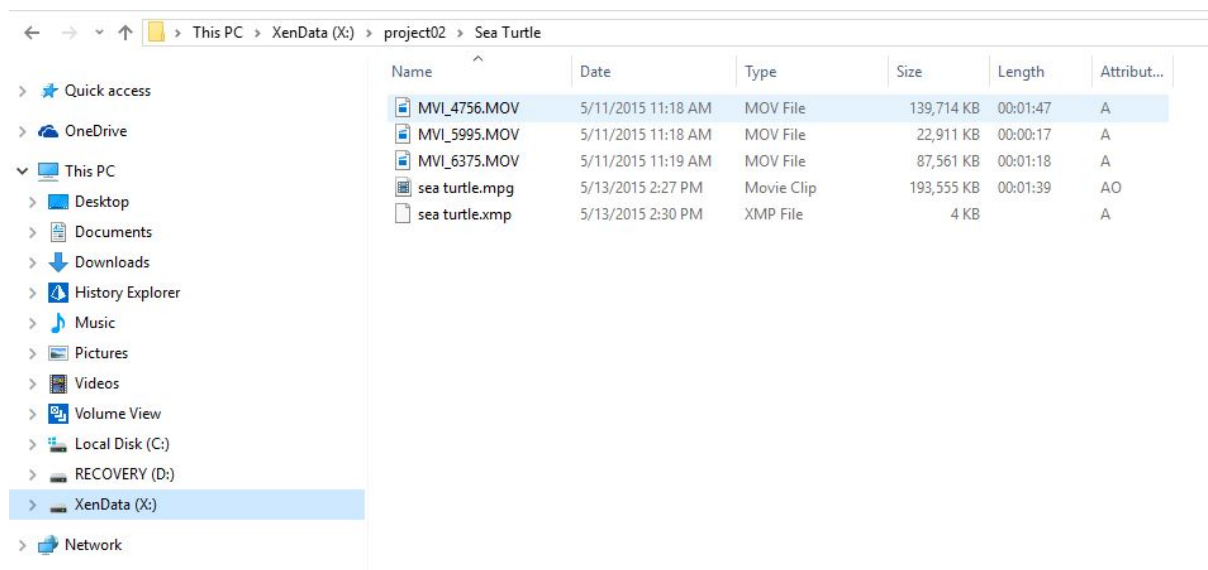
- ❖ If a file is renamed or deleted and then a new file of the same name is created, the system starts again with version 0 of the new file and a new generation is created. The generation number starts at 0 the first time the file is created and increases by one every time a new file is created.

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.

6.4 About Offline Files

When a file has been written to its designated locations, it becomes eligible for flushing from cache. After flushing, the full file is no longer retained on the disk cache and is replaced with a stub file. The stub file has all the same properties as the original except the Microsoft offline attribute is set indicating that the full file is no longer immediately available.

The Windows offline file attribute identifies files that are no longer present on the cache disk. It also increases network timeout periods when a file is being accessed over a network from a Windows client computer.



The image above illustrates how you can use Windows File Explorer to identify offline files in the file system. The file 'sea turtle.mpg' is the only file that is no longer online (i.e. is nearline or offline) as indicated by the offline attribute being displayed. The other four files are stored as full files on the cache disk.

6.5 Handling of Alternate Data Streams

Alternate data streams, also known as 'NTFS streams' and 'named streams', are additional data streams that can be included within a file. Alternate data streams are handled in the following ways:

- ❖ Mac OS/X clients from version 10.6 use alternate data streams when connected to a Windows NTFS share over SMB including the share of a volume managed by the XenData Archive Series software. These alternate data streams contain application-specific file metadata and/or Finder display layout information. The Archive Series software preserves Finder display information on the cache disk but does not write it to the LTO, ODA or object storage account.
- ❖ Windows Internet Explorer adds a stream named 'Zone.Identifier' to files downloaded from the Internet. Windows uses this data for security purposes. The Archive Series software preserves this information on the disk cache but does not attempt to write it to LTO, ODA or object storage.
- ❖ Other types of application-specific alternate data streams will be written to Archival Storage in addition to the disk.

6.6 Supported Network Protocols

You can use CIFS/SMB, NFS and FTP or local file transfers.

You create a file share as you would for a standard Windows logical drive using the standard Microsoft utilities.

6.7 Free Space Reporting

The amount of free space in a storage system is defined as the difference between the total capacity of the system and the amount of space that is used for file data and file system metadata. When Archive Series software is used, definitions of total capacity and free space are as follows:

- ❖ When LTO or ODA storage is employed, the total capacity of the system is defined as the capacity of all volumes known to the system plus the capacity of all blank cartridges. The free space is defined as the total free space on all writable volumes plus the capacity of all blank cartridges.
- ❖ When only Object Storage is employed, the total space is defined by the licensed capacity of the system and the free space is the difference between that value and the amount of data under control by the system.

6.8 File Security

The Archive Series software can be installed within a Windows domain or workgroup. It 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.

When retention of deleted files and old versions of files is enabled, the security model is extended to deleted files and old versions of files. In these cases, the security allocated to prior versions of a file or folder 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.

7. Concepts

The XenData software is easy to administer after understanding a few key concepts, including File Groups, Volumes and Volume Sets.

7.1 About 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, the Archive Series 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.

7.2 About Volumes and Volume Sets

For LTO, the term 'Volume' refers to a complete set of replica tape cartridges which, when up-to-date, all contain the same data. If replication is not enabled, an LTO Volume refers to an individual tape cartridge. For ODA, a Volume is an individual Optical Disc Archive cartridge, for Azure storage accounts, a Volume is a Blob Container, and for Amazon and Wasabi S3 accounts, a Volume is an S3 Bucket.

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 in the Tiered Storage Management Console, the system will automatically add another Volume to extend the Volume Set.

If LTO cartridge 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 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.

One special Volume Set, applicable to LTO and ODA, is the Blank Cartridge Set which 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 Tiered Storage Management Console.

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 XenData control are automatically allocated to File Groups.

7.3 About Volume Catalogs

A Volume Catalog contains an index of the files and folders on the 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.

Some types of Volume may be [Finalized](#) which prevents additional files being written to that Volume. The Finalization process writes the Volume Catalog to the LTO cartridge, a dedicated Azure Container or S3 Bucket.

For LTO and ODA Volumes, when the 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 XenData 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 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.

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.

7.4 About Volume Finalization

Volume finalization is only applicable to TAR formatted tapes, ODA WORM cartridges, Object Storage Containers and S3 Buckets and it prevents additional files being written to that Volume. Finalization is performed automatically when a Volume becomes full and may be initiated manually from the Tiered Storage Management Console. The Finalization process writes the [Volume Catalog](#) to the cartridge, Azure storage account or S3 Bucket.

For LTO and ODA, Finalization 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 XenData Archive Series 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 XenData system to another or when rebuilding a system from the tape cartridges.

For Object Storage, Finalization writes the Volume Catalog to a separate Blob Container or S3 Bucket.

7.5 About Repacking Volumes

Repack is an operation that copies files from one Volume to another, omitting deleted files and old versions of files. It does not change the location of files in the file system i.e. in the file-folder structure.

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). Furthermore, it may be used to move data from one Volume type to another, such as from LTO to Object Storage. 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 Object Set.

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

The repack operation cannot be performed on an archive with only one stand-alone tape or optical drive, unless repacking to Object Storage.

7.6 About Quarantined Objects

If an LTO cartridge, an ODA cartridge, an Azure Blob Container or an S3 Bucket is not usable by the system, it will be identified as a Quarantined Object.

For LTO and ODA, 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 a format that is not compatible with the XenData software; or an error occurred while the system was trying to identify the contents of the cartridge.

7.7 About Pending Write Mode

In normal operation, the Archive Series software writes files to the designated LTO, ODA or object storage immediately after they have been written to the disk cache. However, if the designated storage is not available for any reason, the setting described in [Configuring a Volume Set](#) determines the system's response to an attempt to write files. The response depends on the **Write to disk if no writable Volumes are available** setting. If this is enabled and all writable Volumes in the designated Volume Set become unavailable, the system automatically enters the Pending Write Mode and will accept more data which will be written to the disk cache.

When the system enters the 'Pending Write Mode', it defers writing to the designated LTO, ODA or object storage and continues writing to the disk cache. When a writable Volume becomes available within the Volume Set, the system automatically 'catches up' and writes the pending files to the applicable 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 disk cache. 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.

7.8 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.

When restoring from LTO cartridges, Partial File Restore is enabled in the Archive Series system by using [File Fragmentation](#). File Fragmentation is an optional feature that can be enabled on tape-based systems. It is usually only worthwhile for files that have a size of several tens of

gigabytes or more. When restoring from ODA cartridges or Object Storage, Partial File Restore is supported without need for file fragmentation or any special settings.

Certain application areas, such as Oil and Gas Exploration, generate extremely large files that are bigger than the capacity of the LTO cartridges that are being used. In these applications it is useful to be able to span individual files across multiple cartridges. This can be achieved by enabling [File Fragmentation](#) and a [File Group Advanced Option](#). File spanning is not supported for ODA cartridges.

7.9 Offline File Management

XenData Archive Series software can be configured to provide three tiers of storage hierarchy:

- ❖ **Online** with one instance of a file on the disk cache. In addition, there may be file instances on LTO, ODA or object storage. Files that are read when they are in this state will be restored from the disk.
- ❖ **Nearline** with at least one instance of a file available on LTO, ODA or object storage and not present on the cache disk. Files that are read when they are in this state will be restored from Archival Storage and retained for a predefined time on the disk cache.
- ❖ **Offline** not present on the cache disk and no instance on LTO, ODA or object storage available to the system. In the case of LTO and ODA, it will be unavailable because the cartridge or cartridges containing the file have been exported from the managed libraries or stand-alone drives. In the case of Object Storage, it will be due to loss of connection to the Object Storage account that contains the file.

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 Archive Series software puts a message into the Windows Event Log that identifies which data cartridges, cloud container or bucket 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, container or bucket that contains the file.

7.10 Handling of File Delete and Rename Operations

With LTO and ODA, delete and rename records are written to the Volume that contains the file. This avoids the possibility that delete and rename records written to one cartridge can modify files held on another, making each cartridge self-contained. The Archive Series software does not support deleting or renaming files unless the applicable data cartridge is available for writing, i.e. mounted in a stand-alone drive or in a robotic library, and not write-protected. If there is an attempt to delete or rename a file that is not on an available cartridge, a message is logged in the [Windows Event Log](#) which states that the required cartridge is offline and gives the cartridge

barcode label. When the cartridge is made available, you will be able to perform the rename or delete operation.

If the file is written to a replicated LTO 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 Tiered Storage 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".

8. Administering the System

The main interface for managing the system is the Tiered Storage Management Console which is used to configure all Volume Set and File Group options, including disk cache retention policies. In addition, the Cloud File Gateway uses the Azure Storage Account Configuration and S3 Endpoint Configuration utilities to add and configure Object Storage account access.

8.1 Tiered Storage Management Console

The Tiered Storage Management Console is used to configure all File Group and Volume Set options, manage the operation of Volume Sets and to view diagnostic information about the system. It is a Microsoft Management Console (MMC) snap-in and is illustrated below.




To Start the Tiered Storage Management Console


1. Click the Windows Start icon.
2. Open the XenData program group
3. Click the **XenData System Configuration** entry in the list.

To Navigate the Tiered Storage Management Console

When the console first opens it looks like this:



It shows the logical drive letter under control in the left pane. Click the > symbol to expand the left pane which will then show Configuration, Operation and Diagnostics as shown below.



8.2 Configuring LTO and ODA Storage

When LTO and/or ODA storage hardware is attached and licensed, the Diagnostics section of the Tiered Storage Management Console may be used to configure the following:


- ❖ to disable and enable individual [LTO and ODA libraries and drives](#) which is useful in case of hardware failure; and
- ❖ to [configure Logical Block Protection](#) for LTO hardware. Note that [Logical Block Protection](#) is available for LTO-5 and later generations of drives; it is not available for ODA drives.

8.2.1 LTO Logical Block Protection

[Logical block protection](#) is enable for LTO hardware using the Tiered Storage Management Console.

To Enable Logical Block Protection:

1. Open the Tiered Storage Management Console.
2. Navigate to the **Diagnostics** section.
3. Left-click on the library component.
4. In the right-hand pane, click on the '**Enable Logical Block Protection**' box.




8.3 Configuring Azure Storage Accounts

The Cloud File Gateway uses the Azure Storage Account Configuration utility to add and configure Azure storage account access.


8.3.1 Adding Azure Storage Account Access

1. Launch the Azure Storage Account Configuration utility as follows:
 1. Click the Windows Start icon.
 2. Open the XenData program group
 3. Click the **Azure Storage Account Configuration** entry in the list.


2. Right-click on 'Storage Accounts'; select 'New' and then 'Storage Account'.



3. Enter the name for the storage account (no spaces allowed), then click 'OK'



4. Left-click on the storage account name shown under 'Storage Accounts', choose 'Account Key' or 'Shared Access Signature' depending on the type of access token you will be using, then enter the access token and click 'Apply'.

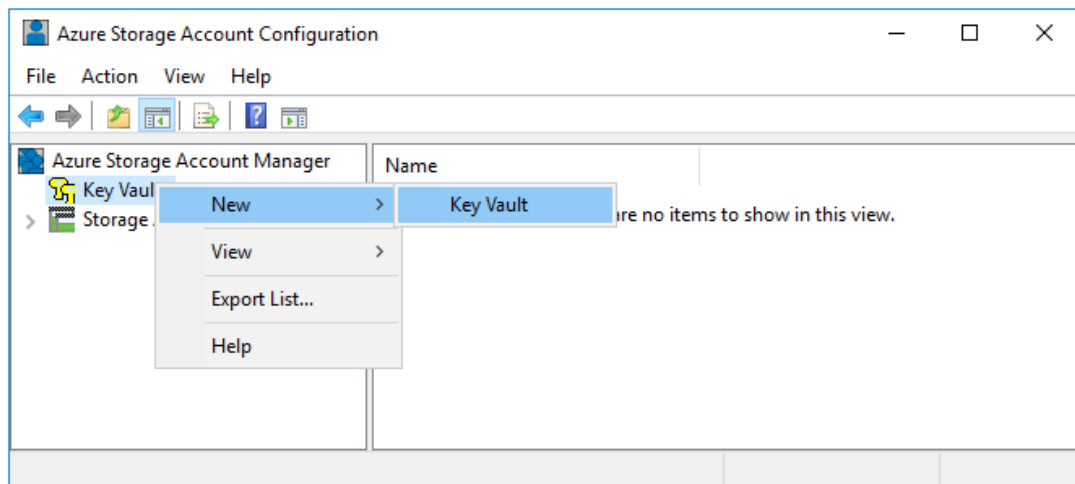


5. Reboot the computer.

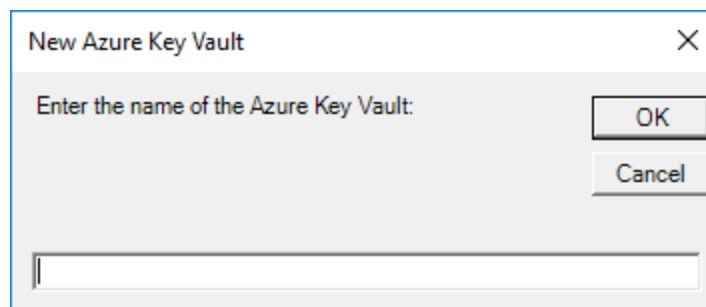
8.3.2 Adding Azure Key Vault Access

An Azure Key Vault can be configured to manage access to one or more Azure Storage Accounts. You can give XenData Archive Series access to the accounts controlled by a key vault by giving it credentials for the key vault.

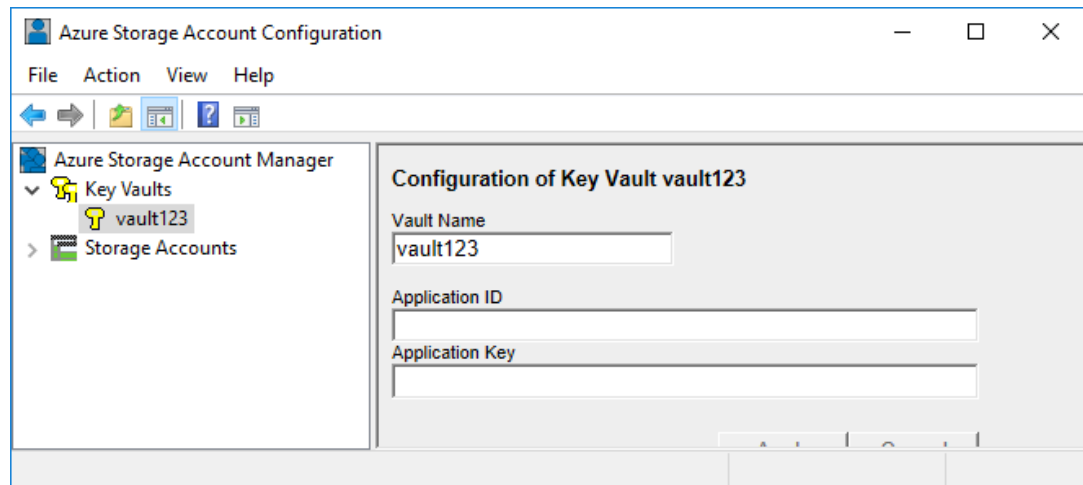
1. Launch the Azure Storage Account Configuration utility as follows:
 1. Click the Windows Start icon.
 2. Open the XenData program group.
 3. Click **Azure Storage Account Configuration** in the list.
2. Right-click on 'Key Vaults'; select 'New' and then 'Key Vault'.



3. Enter the name for the key vault (no spaces allowed), then click 'OK'



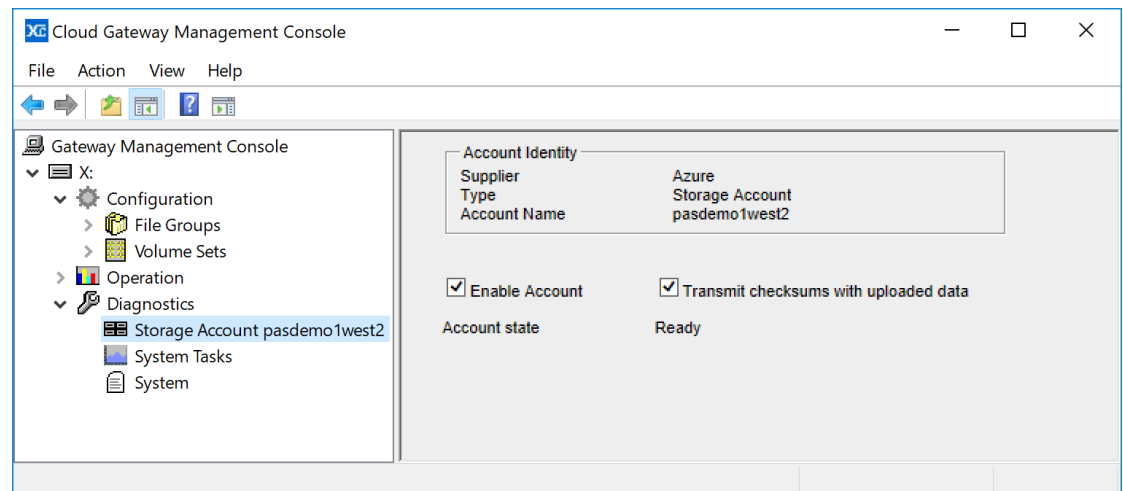
4. Left-click on the key vault name shown under 'Key Vaults', enter the settings for the key vault, then click 'Apply'.



5. Reboot the computer.

8.3.3 Configuring a Storage Account

1. Expand the **Diagnostics** section in the left pane of the Tiered Storage Management Console
2. Click on the Storage Account to be configured



The right-hand pane of the console will show the Account Identity which includes the storage account name. There are two configuration options:


- ❖ **Enable Account.** This must be enabled to access the storage account.
- ❖ **Transmit checksums with uploaded data.** By enabling this option, checksums are transmitted when data is uploaded to the storage account and are used for data verification purposes.

8.4 Configuring Amazon S3 Endpoints


The Cloud File Gateway uses the S3 Endpoint Configuration utility to add and configure Amazon S3 Bucket access.

8.4.1 Adding Amazon S3 Account Access


1. Launch the S3 Endpoint Configuration utility as follows:
 1. Click the Windows Start icon.
 2. Open the XenData program group
 3. Click the **S3 Endpoint Configuration** entry in the list.
2. Right-click on 'S3 Endpoints'; select 'New' and then 'Amazon S3 Endpoint'.



3. Enter the name for the endpoint (no spaces allowed), then click 'OK'




4. Left-click on the endpoint name shown under 'S3 Endpoints', then enter the 'Access Key ID' and 'Secret Key' from your S3 account. Once you've added the keys, select the region you wish your Buckets to be created in from the drop down, and click 'Apply'.



5. Reboot the computer.

8.4.2 Configuring an Amazon S3 Account

1. Expand the **Diagnostics** section in the left pane of the Tiered Storage Management Console
2. Click on the S3 Endpoint to be configured



The right-hand pane of the console will show the Account Identity which includes the endpoint name. There are two configuration options:

- ❖ **Enable Account.** This must be enabled to access the endpoint.

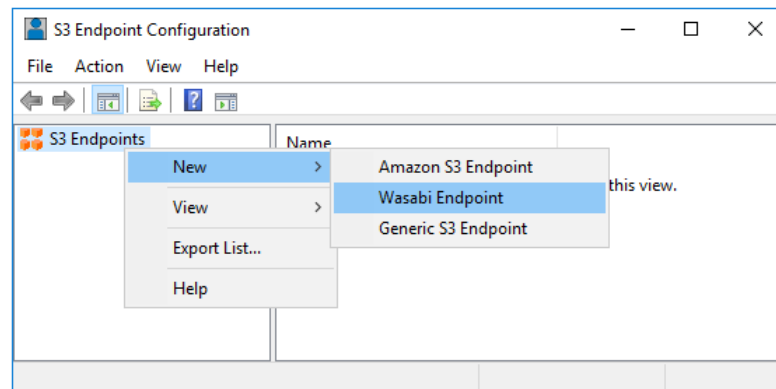
- ❖ **Transmit checksums with uploaded data.** By enabling this option, checksums are transmitted when data is uploaded to the endpoint and are used for data verification purposes.

8.5 Configuring Wasabi S3 Endpoints

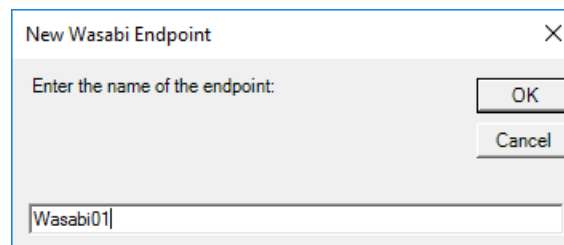
The Cloud File Gateway uses the S3 Endpoint Configuration utility to add and configure Wasabi S3 Bucket access.

8.5.1 Adding Wasabi S3 Account Access

1. Launch the S3 Endpoint Configuration utility as follows:
 1. Click the Windows Start icon.
 2. Open the XenData program group
 3. Click the **S3 Endpoint Configuration** entry in the list.
2. Right-click on 'S3 Endpoints'; select 'New' and then 'Wasabi Endpoint'.

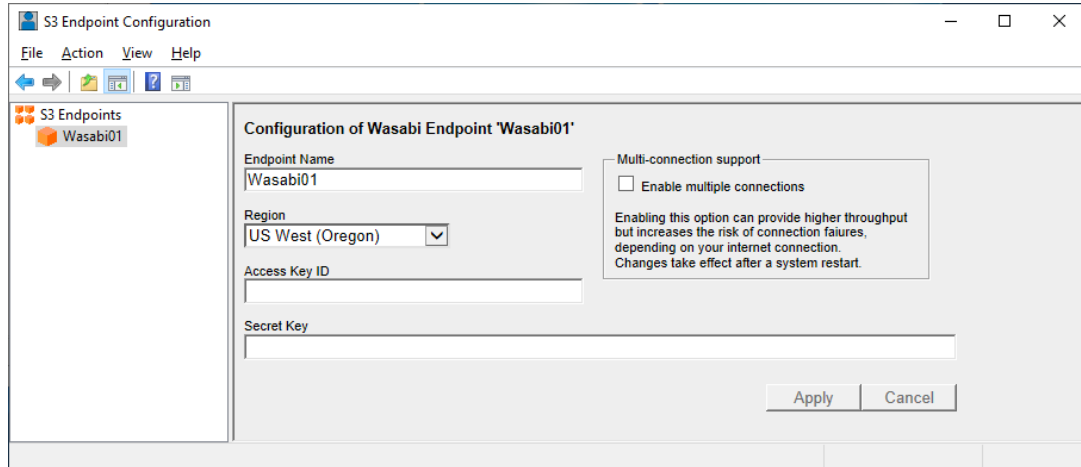


3. Enter the name for the endpoint (no spaces allowed), then click 'OK'



4. Left-click on the endpoint name shown under 'S3 Endpoints', then enter the 'Access Key ID' and 'Secret Key' from your S3 account. Once you've added the keys, select the region you wish your Buckets to be created in from the drop down, and click 'Apply'. You can optionally enable 'Multi-Connection support', this allows the system to send multiple

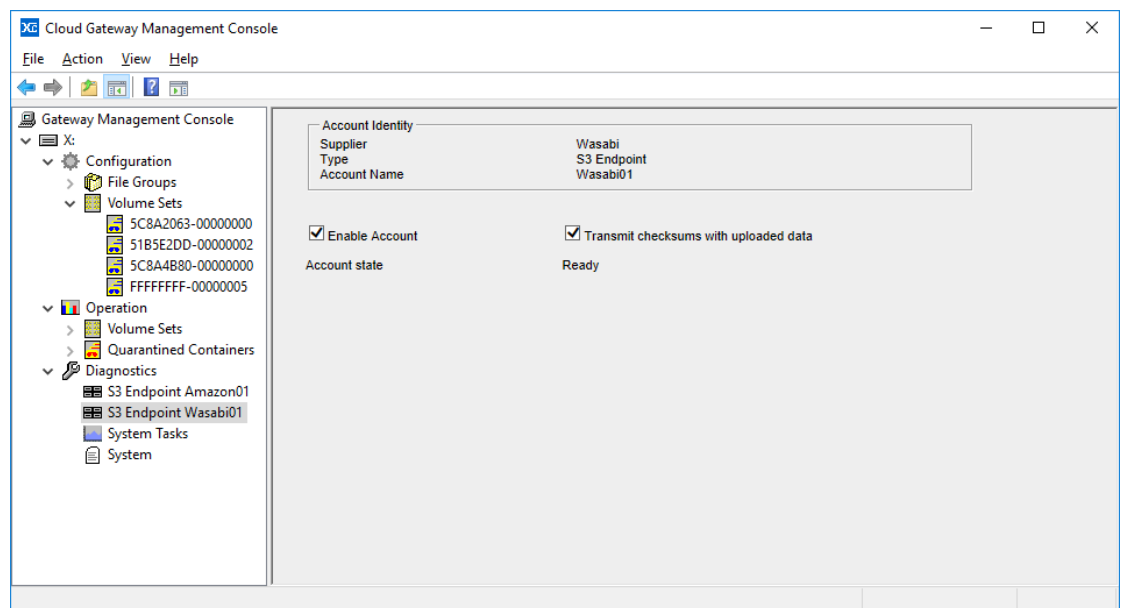
streams of data to the S3 bucket for increased performance. This option has a higher risk of write failures on some internet connections, so should be tested on your connection before production use.



5. Reboot the computer.

8.5.2 Configuring a Wasabi S3 Account

1. Expand the **Diagnostics** section in the left pane of the Tiered Storage Management Console
2. Click on the S3 Endpoint to be configured



The right-hand pane of the console will show the Account Identity which includes the endpoint name. There are two configuration options:


- ❖ **Enable Account.** This must be enabled to access the endpoint.
- ❖ **Transmit checksums with uploaded data.** By enabling this option, checksums are transmitted when data is uploaded to the endpoint and are used for data verification purposes.

8.6 Configuring Generic S3 Endpoints


The Cloud File Gateway uses the S3 Endpoint Configuration utility to add and configure Generic S3 Bucket access.

8.6.1 Adding Generic S3 Account Access


1. Launch the S3 Endpoint Configuration utility as follows:
 1. Click the Windows Start icon.
 2. Open the XenData program group
 3. Click the **S3 Endpoint Configuration** entry in the list.
2. Right-click on 'S3 Endpoints'; select 'New' and then 'Generic Endpoint'.



3. Enter the name for the endpoint (no spaces allowed), then click 'OK'




4. Left-click on the endpoint name shown under 'S3 Endpoints', then enter the 'Endpoint URL', 'Region', 'Access Key ID' and 'Secret Key' from your S3 account. Once you've entered all your account information click 'Apply'. You can optionally enable 'Multi-Connection support', this allows the system to send multiple streams of data to the S3 bucket for increased performance. This option has a higher risk of write failures on some internet connections, so should be tested on your connection before production use.



5. Reboot the computer.

8.6.2 Configuring a Generic S3 Account

1. Expand the **Diagnostics** section in the left pane of the Tiered Storage Management Console
2. Click on the S3 Endpoint to be configured



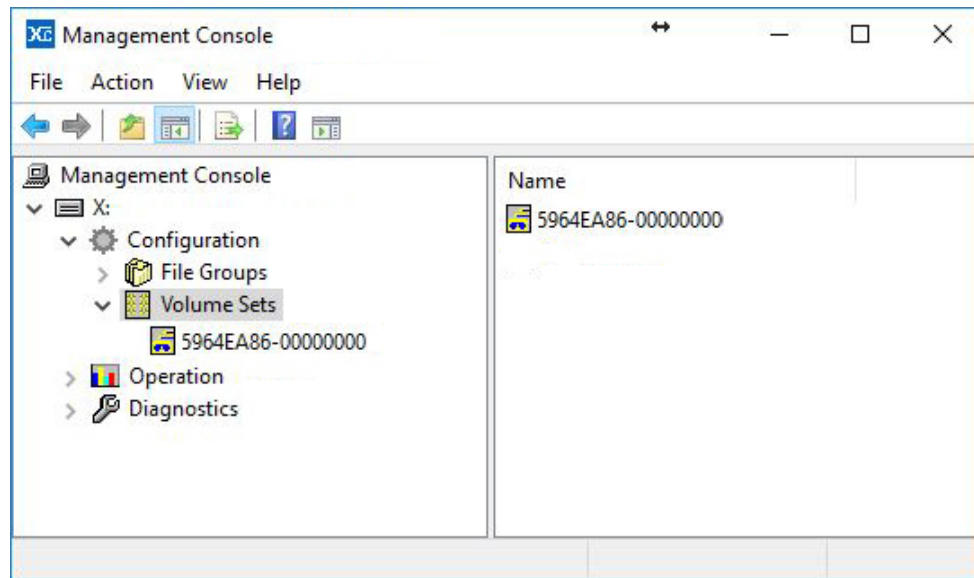
The right-hand pane of the console will show the Account Identity which includes the endpoint name. There are two configuration options:

- ❖ **Enable Account.** This must be enabled to access the endpoint.
- ❖ **Transmit checksums with uploaded data.** By enabling this option, checksums are transmitted when data is uploaded to the endpoint and are used for data verification purposes.

8.7 Volume Sets

A Volume Set is a set of one or more Volumes that store files from designated File Groups. A Volume Set expands dynamically, adding Volumes as needed.

For a new installation of the Archive Series software, an initial Volume Set is automatically created ready for configuration, as illustrated below.



8.7.1 Creating a New Volume Set

1. Expand the **Configuration** section in the left pane of the Tiered Storage Management Console
2. Right click on Volume Sets
3. Click on New --> Volume Set

The new Volume Set is now ready to be configured, as described in [Configuring a Volume Set for Object Storage](#) and [Configuring a Volume Set for LTO or ODA](#).

8.7.2 Renaming a Volume Set

1. Expand the **Configuration** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Right click on the Volume Set to be renamed
4. Click on **Rename**
5. Rename the Volume Set and press **Enter**


8.7.3 Deleting a Volume Set

1. Expand the **Configuration** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Right click on the Volume Set to be deleted
4. Click on **Delete**

Note that a Volume Set that has Volumes allocated to it cannot be deleted. The Volumes within the Volume Set must first be deleted as described in [Deleting a Volume](#).

8.7.4 Configuring a Volume Set for LTO or ODA

1. Expand the **Configuration** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Click on the Volume Set to be configured



Configure the selected Volume Set in the right hand pane of the console as follows:

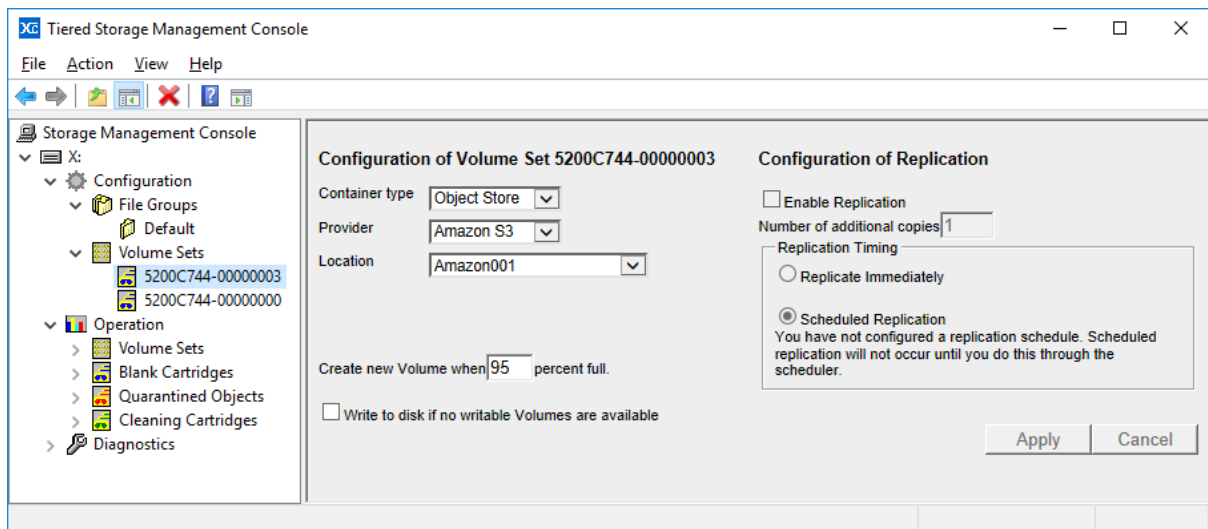
1. Select Tape or Optical for the Cartridge type, as appropriate for your hardware.
2. For tape cartridges, select the File system, either TAR or LTFS.
3. Select WORM or Rewritable data cartridges using the two radio buttons.
4. The Compressed check box enables data compression if the hardware supports 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.

After having configured these fields, click Apply. Note that if you are configuring a new Volume Set, a first Volume must be added before it is ready for use. This operation is described in [Adding a Volume](#).

8.7.5 Configuring a Volume Set for Object Storage

1. Expand the **Configuration** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Click on the Volume Set to be configured



In the right hand pane of the console, the options shown in the Container type field will be determined by how the system is licensed. Select Object Store as the Container Type. The File System will show a list of supported providers, depending on what type of Object Storage you are using. The location drop down allows you to specify which Object Store you wish that volume set to write to, if you have multiple of the same type. If you only have a single Object Store of that type, you won't need to set this, and it can be left blank. There are two fields that can then be configured:

- ❖ **Create new Volume when x percent full.** This determines the percentage full of the current Volume at which a new Volume is automatically added. A Volume becomes full when it contains 1 million Objects and this setting has a default value of 95% which represents 950,000 Blobs. To prevent the automatic creation of a new Volume, set this value to 100%.
- ❖ **Write to disk if no writable Volumes are available.** This option determines the system's behavior if all Volumes in the Volume Set become unavailable. The Volumes may become unavailable, for example, if an Internet connection is lost. If this option has been enabled and all Volumes in a Volume Set become 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. This is described further in [About Pending Write Mode](#).

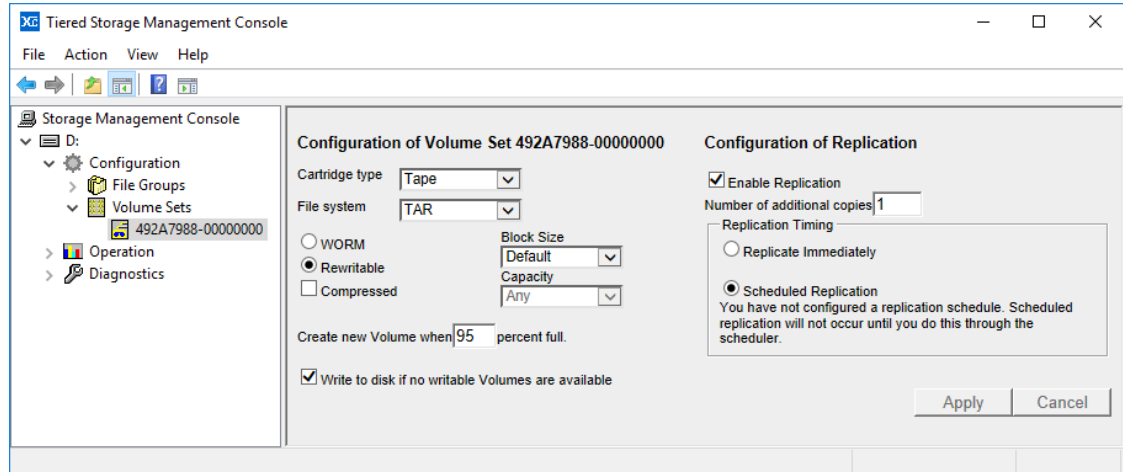
After having configured these fields, click Apply. Note that if you are configuring a new Volume Set, a Volume must first be added before it is ready for use. This operation is described in [Adding a Volume](#).

8.7.6 Configuring Replication for an LTO Volume Set


Replication settings for an LTO 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 an LTO system that has no library and only a single stand-alone tape drive.

Replication is defined for a Volume Set as follows:

1. Expand the **Configuration** section in the left pane of the Tiered Storage Management Console.
2. Expand the **Volume Sets** section.
3. Select the required Volume Set to reveal the "Configuration of Volume Set" panel in the right pane of the window.
4. Check the **Enable Replication** box.
5. Enter the required number of replicas in the **Number of additional copies** box.
6. If the replication is to be performed immediately, then check the **Replicate Immediately** box.



1. If the replication is to be performed at a future date or time, then check the **Scheduled Replication** box.



2. Click on **Apply**.

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

8.7.7 Adding a Volume

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Right click on the Volume Set to which the Volume is to be added

4. Click on **Add Volume**

This creates a new Volume allocated to the selected Volume Set.

Note: the system automatically adds a Volume to a Volume Set when the current Volume reaches the percentage full that is defined in [Configuring a Volume Set for LTO or ODA](#) and [Configuring a Volume Set](#). However, the first Volume in a Volume Set must be created manually as described here.

8.7.8 Displaying Information about a Cartridge

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Expand the required Volume Set to show the Volumes that are allocated to it.
4. Click on a cartridge to show the information pane

The screenshot displays the Tiered Storage Management Console interface. The left pane shows a tree view under 'Storage Management Console' with 'Operation' expanded to 'Volume Sets', and a specific cartridge 'Barcode:04347AL6' selected. The right pane shows the 'Information' tab for this cartridge, listing various attributes and a pie chart representing storage usage.

Information	Identity
Cartridge Identity	Barcode:04347AL6
Location	Library 0 slot 1
State	Primary
Mount Cycles	28
Last alert	None
Last verified	Never
Verify Status	None
0.0% used (2328GB free out of 2328GB)	
Overhead: 95.8% (0GB)	
Info	<input type="button" value="Update"/>


The pie chart below the table shows storage 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.



8.7.9 Adding User Defined Information for a Cartridge

To add user-defined information about a cartridge:

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Expand the required Volume Set to show the Volumes that are allocated to it.
4. Click on a cartridge to show the information pane
5. Enter the user defined information in the info field, then click **Update**




8.7.10 Verifying Cartridges

The verify function performs a block level check of data written to an LTO tape cartridge, and is useful for checking the integrity of the LTO Tape cartridge and LTO Tape Drive hardware. This is not supported for ODA hardware.

To Verify the Data on a Cartridge

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Expand the required Volume Set to show the Volumes that are allocated to it.
4. Right click on the cartridge and select **Verify**.



8.7.11 Reformatting Rewritable LTO or ODA Cartridges

Reformatting a rewritable LTO or ODA 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: "Offline/archived (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 Tiered Storage Management Console.

To Reformat a Rewritable Cartridge:

1. Expand the Operation section in the left pane of the Tiered Storage Management Console
2. Expand the Volume Sets section
3. Expand the required Volume Set to show the Volumes that are allocated to it.
4. Click on a cartridge, then click Reformat in the right pane, or, right-click on the Volume and select Reformat.
5. 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.

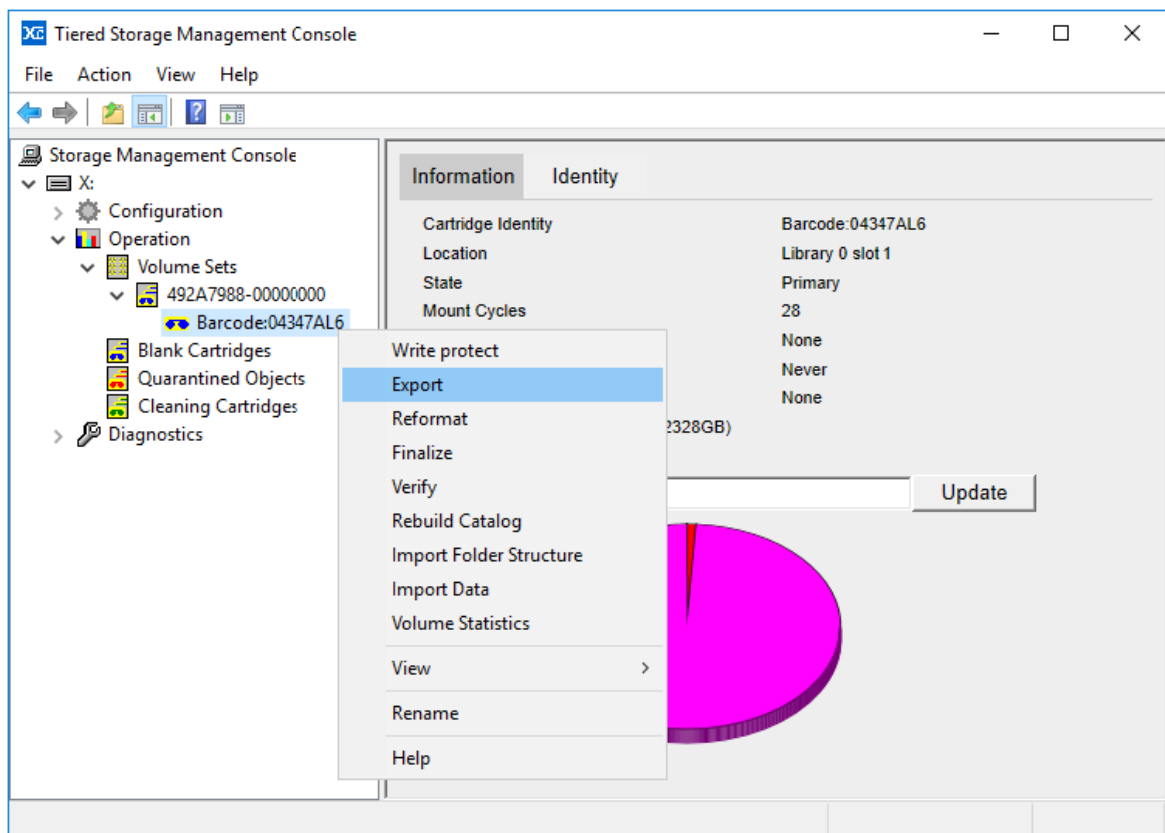


8.7.12 Exporting Cartridges from an LTO or ODA Library

Cartridges should be exported from a robotic library using the Tiered Storage Management Console as described below and not by using the front panel of the library, or library web interface.

To Export a Data Cartridge

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Expand the required Volume Set to show the Volumes that are allocated to it.
4. Click on the Cartridge to be exported, then click **Export** in the right pane of the console, or right-click the cartridge 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 or optical drive, a cartridge may be exported (ejected) either by using the procedure described above or by pressing the drive eject button.

8.7.13 Scanning for Object Storage Containers Created by Other Systems

This operation will identify any Containers or Buckets that have been created in the accessible Object Storage Accounts by another system including by another XenData Cloud File Gateway or by a third party application. Scanning is performed as follows:

1. Expand the Diagnostics section in the left pane of the Tiered Storage Management Console
2. Right click on the Object Storage Account
3. Click on Scan for new Containers



In the right pane of the console, the options shown in the Container type field will be determined by how the system is licensed. Select Object Store as the Container Type. The File System will be shown as either Amazon S3, Azure or Wasabi, depending on what type of object storage you are using. The Location drop down allows you to specify which Object Store you wish that Volume Set to write to, if you have multiple Locations. If you only have a single Location for that File System, you will not need to set this, and it can be left blank. There are two fields that can then be configured:

- ❖ Create new Volume when x percent full. This determines the percentage full of the current Volume at which a new Volume is automatically added. A Volume becomes full when it contains 1 million objects and this setting has a default value of 95% which represents 950,000 objects. To prevent the automatic creation of a new Volume, set this value to 100%.
- ❖ Write to disk if no writable Volumes are available. This option determines the system's behavior if all Volumes in the Volume Set become unavailable. The Volumes may become unavailable, for example, if an Internet connection is lost. If this option has been enabled and

all Volumes in a Volume Set become 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. This is described further in About Pending Write Mode.

After having configured these fields, click Apply. Note that if you are configuring a new Volume Set, a Volume must first be added before it is ready for use. This operation is described in [Adding a Volume](#).

Note that when the Cloud File Gateway software starts up, for example when the machine reboots, it scans the object storage accounts available and identifies any new Containers or Buckets and adds them as Volumes in the console, avoiding the need to use the scan operation described here.

8.7.14 Deleting an Object Storage Container

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Right click on the Volume to be deleted
4. Click on **Delete Container**
5. Click **OK** to delete the Container or Bucket

Note: applicable to systems with multiple Cloud File Gateway instances: only Volumes created by this instance of the Cloud File Gateway can be deleted.

8.7.15 Rebuilding Volume Contents Catalogs

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Right click on the Volume to be rebuilt
4. Click on **Rebuild Catalog**

This operation is useful when importing files written to Object Storage written by another system. Its operation will update the Volume Catalog for the selected Volume. It will not change the file contents.

For more information please refer to [Importing Files Written to Object Storage by Another System](#).

8.7.16 Import Folder Structure

The Import Folder Structure operation is applied to an individual Volume or to all the Volumes in a Volume Set and it updates the file-folder interface with the file structure defined in the Volume Catalogs for the applicable Volumes. After the operation is complete, files will appear in the file-folder interface as 'offline', which means they are present, but will need to be restored from the storage medium before access.

Perform the Import Folder Structure operation for a selected Volume as follows:

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Expand the required Volume Set
4. Right click on the required Volume
5. Click on **Import Folder Structure**

Perform the Import Folder Structure operation for all the Volumes in a selected Volume Set as follows:

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Right click on the required Volume Set
4. Click on **Import Folder Structure**

8.7.17 Import Data

The Import Data operation is applied to an individual Volume or to all the Volumes in a Volume Set. It updates the file-folder interface with the file structure defined in the Volume catalogs for the Volumes and it also selectively stores file instances on the disk cache in accordance with the Disk Retention Rules for written files.

Perform the Import Data operation for a selected Volume as follows:

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Expand the required Volume Set
4. Right click on the required Volume
5. Click on **Import Data**

Perform the Import Data operation for all the Volumes in a selected Volume Set as follows:

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Right click on the required Volume Set
4. Click on **Import Data**

8.7.18 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. For rewritable LTO and ODA cartridges, 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 storage type to another. For example, repack can be used to migrate files from LTO to Azure Blob Storage or from an older generation LTO cartridge to the latest generation.

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 Object 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.

LTO and ODA Hardware Requirements

The repack operation cannot be performed on an archive without a library and only one standalone drive. However, it can be performed on an archive system with a robotic library having only one drive but this might be a very slow operation.

To Repack a Volume


1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Expand the required Volume Set to show the Volumes that are allocated to it.
4. Right-click on the Volume to be repacked, and ensure that it is not writable. If it is writable, then write-protect it.
5. Select **Repack**.



8.7.19 Cancel Volume Repack

The Repack operation 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.

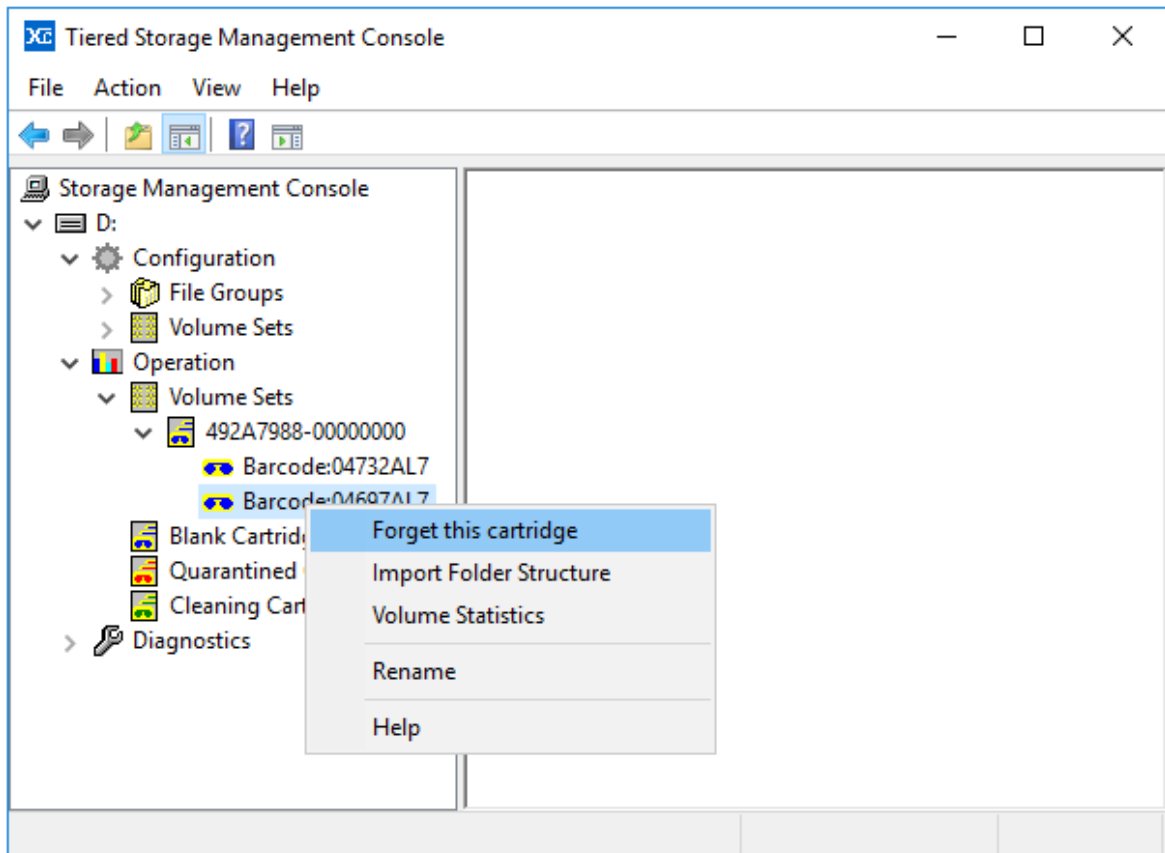


8.7.20 Removing Information about a Cartridge from the System

If a cartridge 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. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Expand the required Volume Set
4. Right click on the required Volume
5. Click on **Forget this cartridge**




8.7.21 Replacing a Missing Replica Cartridge

If Replication is enabled for an LTO Volume Set then files are automatically written 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 LTO cartridge](#).

To Add a Missing Replica

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Expand the required Volume Set
4. Right click on the required Volume
5. Click on **Add missing replica**



8.7.22 The Blank Cartridge Set

The Blank Cartridge Set is a special Volume Set that contains data cartridges that have been imported into a robotic LTO or ODA 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 user. Cartridges in the Blank Cartridge Set are allocated to an operational Volume Set either manually by the user or automatically as defined by the [Volume Set configuration](#) setting.

8.7.23 Cleaning LTO Drives

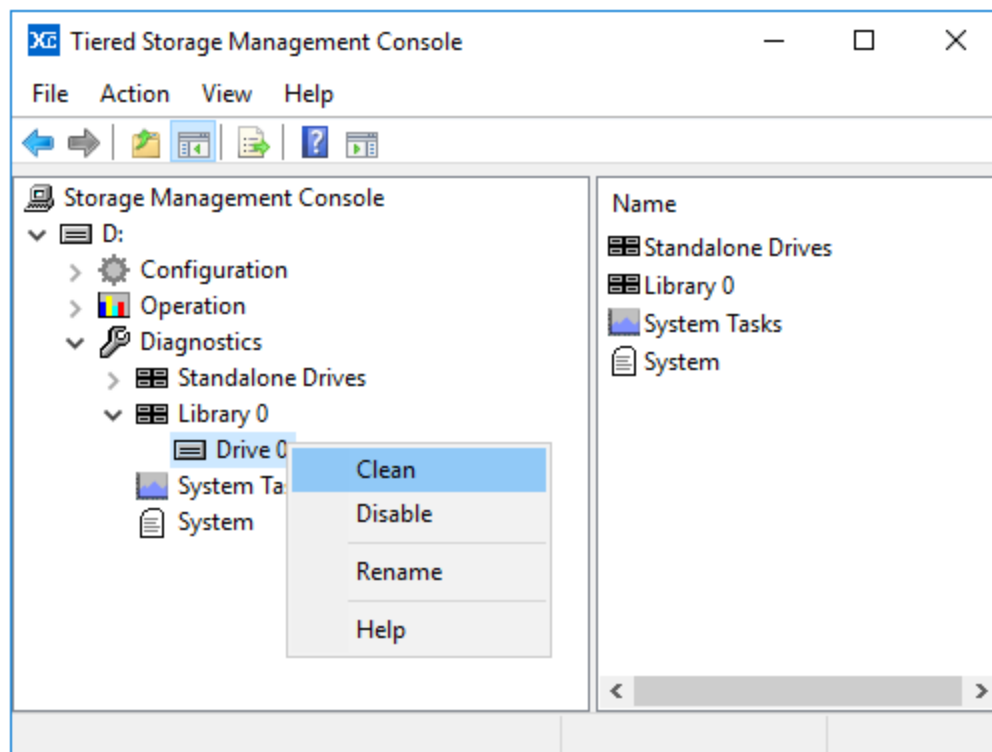
LTO tape drives require periodic cleaning which is performed by inserting a cleaning cartridge in the drive. Normally, a drive will issue a request for cleaning at the appropriate time and in the case of drives within robotic libraries, the Archive Series software responds to these requests by inserting 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 Tiered Storage Management Console identifying that the drive requires cleaning.

If required, you can manually clean a drive within a robotic library as follows:

To Manually Clean an LTO Drive within a Library


1. Expand the **Diagnostics** section in the left pane of the Tiered Storage Management Console
2. Expand the **Library** section and find the appropriate drive.
3. Right-click on the drive and select **Clean**.



8.7.24 Displaying Information about Cleaning Cartridges

XenData Archive Series software recognizes LTO cleaning cartridges and displays information about them in the Operation section of the Tiered Storage Management Console as illustrated below.


Cleaning cartridges are used to clean LTO drives when necessary. The Archive Series software 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.



8.7.25 Quarantined Object Set

Quarantined Objects is a special Volume Set that contains Object Storage Containers or 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. This may be because the contents have been repacked, because a data 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 Object Storage Container or data cartridge.

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

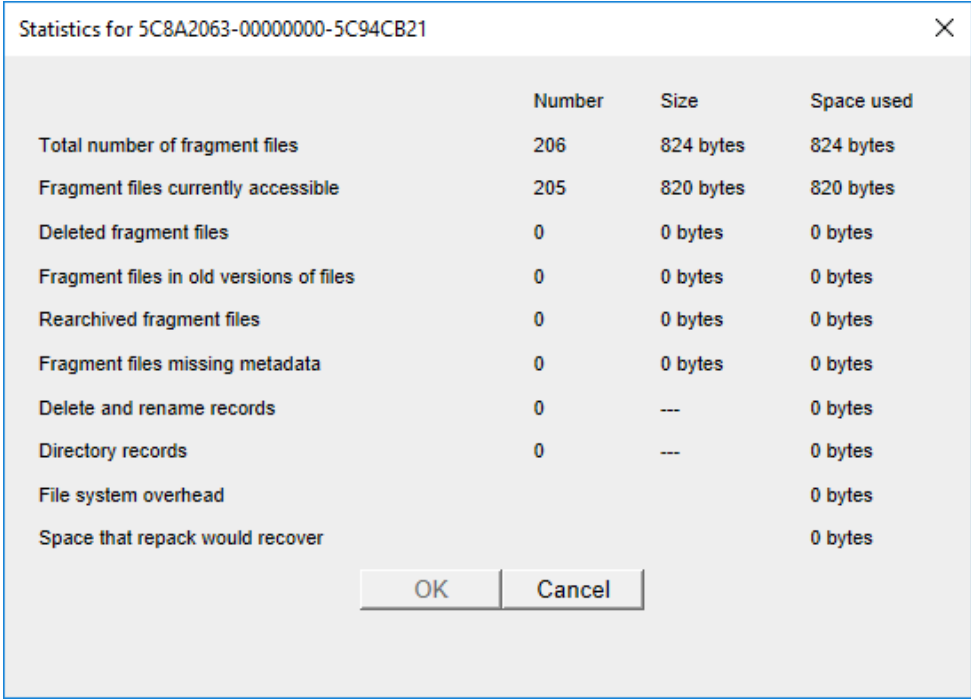


8.7.26 Obtaining Volume Statistics

Volume Statistics displays relevant information about an individual Volume.

To obtain statistics for a Volume:

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Right click on the Volume to be selected
4. Click on **Volume Statistics**



	Number	Size	Space used
Total number of fragment files	206	824 bytes	824 bytes
Fragment files currently accessible	205	820 bytes	820 bytes
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	0	---	0 bytes
File system overhead			0 bytes
Space that repack would recover			0 bytes

8.7.27 Write Protecting a Volume

There may be circumstances when you want to stop the system from writing data to a particular Volume before it becomes full. This can be achieved by write protecting the Volume. If all the Volumes in a Volume Set are full, finalized or write-protected, you will have to add a new Volume before more data can be written to the Volume Set. This is described in [Adding a Volume](#).

To write protect a Volume:

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console

2. Expand the **Volume Sets** section
3. Right click on the Volume to be write protected
4. Click on **Write protect**

Note applicable to systems with multiple Cloud File Gateway instances: only Volumes created by this instance of the Cloud File Gateway can be write protected.

8.7.28 Finalizing Volumes

Volume finalization prevents additional files being written to the Volume. Finalization occurs automatically when a Volume becomes full and may be performed manually as described here. The Finalization process writes the [Volume Catalog](#) to a separate Object Storage Container or Bucket. To Finalize a Volume:

1. Expand the **Operation** section in the left pane of the Tiered Storage Management Console
2. Expand the **Volume Sets** section
3. Right click on the Volume to be Finalized
4. Click on **Finalize**
5. Click **OK** to Finalize the Volume

Note applicable to systems with multiple Cloud File Gateway instances: only Volumes created by this instance of the Cloud File Gateway can be Finalized.

8.8 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. Files are assigned to a File Group on the basis of their name and path.

After initial installation of the Archive Series 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, as described in [Creating a New File Group](#).

8.8.1 Creating a New File Group

To create a new File Group:

1. Expand the **Configuration** section in the left pane of the Tiered Storage Management Console
2. Right click on **File Groups**
3. Click New File Group
4. Enter a name for the new File Group

5. Click **OK**

It should then be edited as described in [Allocating Files to a File Group](#), [Selecting a Volume Set for a File Group](#), [Selecting Disk Retention Rules](#) and [File Group Advanced Options](#).

8.8.2 Renaming a File Group

To rename a File Group, as displayed in the Tiered Storage Management Console:

1. Expand the **Configuration** section in the left pane of the Tiered Storage Management Console
2. Expand the **File Groups** section
3. Right click on the File Group to be renamed
4. Click **Rename**
5. Enter the new name for the File Group

8.8.3 Changing the Order of File Groups

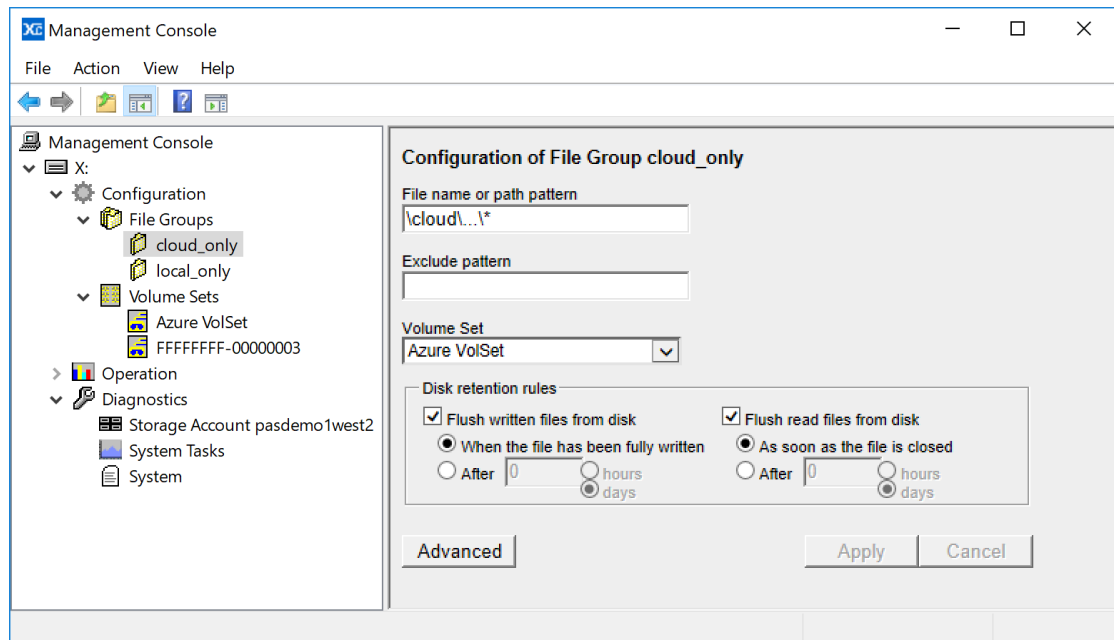
The order of File Groups in the Tiered Storage Management Console is important because an individual file can be allocated to only one File Group and the allocation rules are applied in the order that the File Groups appear in the left pane of the console with files being allocated to the uppermost applicable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

1. Expand the **Configuration** section in the left pane of the Tiered Storage Management Console
2. Expand the **File Groups** section
3. Right click on a File Group to move it up or down
4. Click either **Move Up** or **Move Down**

8.8.4 Allocating Files to a File Group

Files are allocated to File Groups based on their folder name, file name, extension or a combination of these. To allocate files to a File Group:

1. Expand the **Configuration** section in the left pane of the Tiered Storage Management Console
2. Expand the **File Groups** section
3. Click on a File Group to display configuration options in the right pane



4. Update the File name or path pattern box with text to select the required files using the conventions described below
5. If required, update the Exclude pattern box using the conventions described below
6. Click **Apply**

Standard file name and wild card conventions (such as "*" and "?") may be used within the pattern match. As an extension to normal pattern matching syntax, the special folder wild card '...' can be used to match all sub-folders. The system supports multiple patterns per File Group, separated by semicolons. Some example file name or path patterns are:

- ❖ *.mov selects files with the extension .mov for the File Group.
- ❖ abc???.mov selects files that start with abc, have the extension .mov 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\...*.mov selects files with the extension .mov that are in the folder \Images or any of its sub-folders.

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

Note that if there is no matching File Group for a file, the system blocks opening or creation of the file and returns an error to the application that tried to use the file.

8.8.5 Examples of Allocating Files to a File Group

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 logical drive letter managed by the XenData Archive Series software, 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 with 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.

8.8.6 Selecting Storage Options for a File Group

To Select Storage Options for a File Group

1. Expand the **Configuration** and **File Groups** section in the left pane of the Tiered Storage Management Console.
2. Navigate to the File group.
3. Determine whether files in the File Group are to be saved to a Volume Set. If so, select the required Volume Set from the **Volume Set** drop-down menu. Otherwise, select **None** from the **Volume Set** drop-down menu.
4. In the case of saving to an LTO Volume Set, determine whether or not [file fragmentation](#) is required and enable if appropriate. If you enable file fragmentation, the fragment size must be defined. Recommended fragment sizes depend on the application and LTO drive transfer rates but it will typically be 10 GB or larger.
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](#).
6. If appropriate, click on the [Advanced](#) button near the bottom of the screen and make additional selections.
7. Click **Apply**




8.8.7 Selecting a Volume Set for a File Group

1. Expand the **Configuration** section in the left pane of the Tiered Storage Management Console
2. Expand the **File Groups** section
3. Click on the applicable File Group to display configuration options in the right pane

4. Select the required Volume Set from the drop-down options in the Volume Set box
5. Click **Apply**

Note that if the files allocated to the File Group are to be saved only on the cache disk, select **None** as the Volume Set option, as illustrated below.



8.8.8 Selecting File Fragmentation

[File fragmentation](#) is an option applicable to File Groups that write to LTO Volume Sets. It should be enabled to support partial file restores (PFR) from LTO cartridges and/or to support spanning of very large files across LTO cartridges.

Enable file fragmentation as described in [Selecting Storage Options for a File Group](#).

8.8.9 Selecting Disk Retention Rules

You can configure the system such that, after a file has been securely written to a Volume Set, the instance stored on disk will be flushed to release the disk space occupied by the file. Flush functionality is enabled by configuring the Disk retention rules in the File Group configuration options. The Disk retention rules are only available for File Groups where the files are stored on a Volume Set.

Characteristics of flushed files are as follows:

- ❖ Flushing from the cache disk 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 from LTO, ODA or object storage by simply reading the file.

To configure disk retention rules:

1. Expand the **Configuration** section in the left pane of the Tiered Storage Management Console
2. Expand the **File Groups** section
3. Click on the applicable File Group to display configuration options in the right pane
4. Make settings in the Disk retention rules box as described below.
5. Click **Apply**

Examples of common disk retention rules settings are given below:

- ❖ **Files are retained on disk indefinitely.** Deselect both **Flush written files from disk** and **Flush read files from disk** as shown below.

Disk retention rules

Flush written files from disk Flush read files from disk

When the file has been fully written As soon as the file is closed

After 0 hours days After 0 hours days

- ❖ **Files are flushed immediately after writing to a Volume Set and remain flushed after reading.** Select **Flush written files from disk** and **When the file has been fully written**. Also select **Flush read files from disk** and **As soon as the file is closed** as shown below.

Disk retention rules

Flush written files from disk Flush read files from disk

When the file has been fully written As soon as the file is closed

After 0 hours days After 0 hours days

- ❖ **Files are flushed a preset length of time after being writing or last read.** Select **Flush written files from disk** and **After**, choose a number of hours or days. Also select **Flush read files from disk** and **After** the same number of chosen hours or days. With these options selected, files are retained on cache disk for the defined length of time after

they were written or last read. This is illustrated below with options selected to keep files on disk for 60 days after first written or last read.

- ❖ **Files are flushed a preset length of time after writing or immediately after being read.** Select **Flush written files from disk** and **After**, choose a number of hours or days. Also deselect **Flush read files from disk**. With these options selected, files are retained on disk for the defined length of time after they were written or they are flushed immediately after being read for the first time, which ever happens sooner.
- ❖ **Files are retained on disk until they are first read.** Deselect **Flush written files from disk** and select **Flush read files from disk**. Also select **As soon as the file is closed**. With these options selected, files are retained on disk until they are first read after which they are flushed.

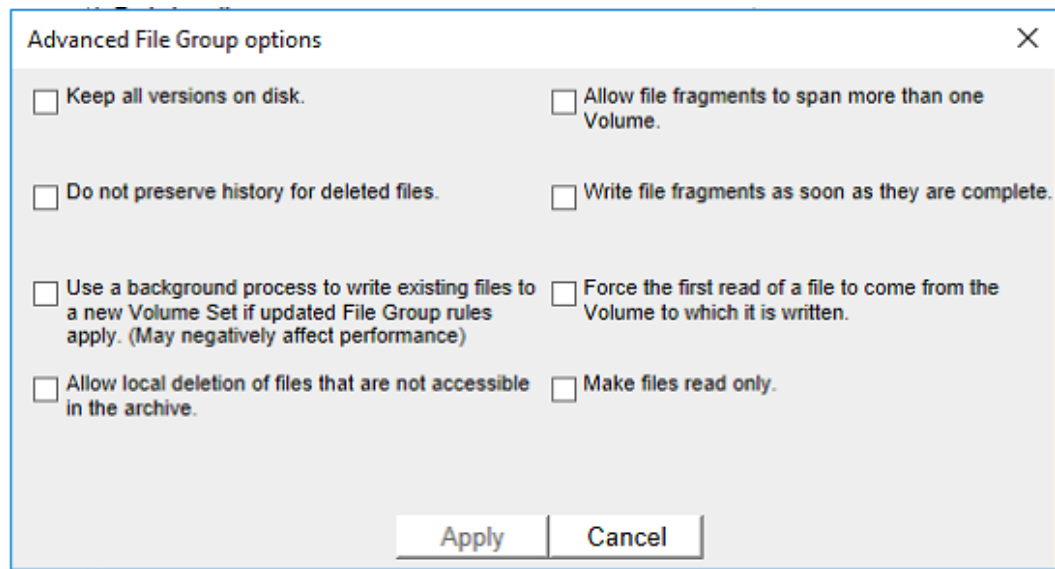
8.8.10 Changing Disk Retention Rules

You 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 disk, you 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.

8.8.11 File Group Advanced Options

To configure advanced File Group options:

1. Expand the **Configuration** section in the left pane of the Tiered Storage Management Console
2. Expand the **File Groups** section
3. Click on the applicable File Group to display configuration options in the right pane
4. Click **Advanced** which will display the advanced File Group options window, as illustrated below
5. Configure required advanced options
6. Click **Apply**



A description of the available options is given below:

- ❖ **Keep all versions on disk.** Default behavior for the system is to keep only the latest version of a file on disk. Selecting **Keep all versions on disk** changes this behavior so that all old versions of a file are retained on disk.
- ❖ **Do not preserve history for deleted files.** In normal operation, the software maintains version history of files under its control. Maintaining a file's history consumes space on 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 subsequently deleted files.
- ❖ **Use a background process to write existing files to a new Volume Set if updated File Group rules apply.** This is a useful option if a File Group setting initially only stored files on disk due to the selected Volume Set having been configured to **None** and then the Volume Set setting was changed to save files to LTO, ODA or object storage. Note that this setting will likely impact performance and should not be enabled permanently.
- ❖ **Allow local deletion of files that are not accessible in the archive.** This option allows the deletion of files from the XenData cache drive, without attempting to delete them from the archive medium (LTO, ODA, Object Storage Service). This allows the deletion of files even if the archive medium is inaccessible.
- ❖ **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 Volumes. When this option is not enabled, all file fragments for a particular version of a file will be written to the same Volume.

- ❖ **Write file fragments as soon as they are complete.** In normal operation, the software writes files to the designated Volume Set after the whole file has been written to disk and the file has been closed. Sometimes there is a requirement to write data to the Volume Set 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 advanced option. File fragmentation must be enabled for this option and normally file fragmentation is applicable only when writing to LTO Volume Sets.
- ❖ **Force the first read of a file after it is written to come from the Volume to which it is written.** Some applications employ a read-after-write check to verify the integrity of data written. 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 system, this will usually be the cache disk (or even an intermediate RAM cache). This option forces data to be read from the LTO, ODA or object storage, even if it is available from cache disk, thereby allowing applications to verify the integrity of data written to LTO, ODA or object storage. **Note** that only the first read is forced to come from the LTO, ODA or object storage; subsequent reads will be satisfied from the cache disk if possible.
- ❖ **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.

8.9 Multi-Site Sync Service

8.9.1 Adding the Multi-Site Sync Service

The [Multi-Site Sync](#) functionality is provided by XenData as a service offering with subscription licensing. The service uses Azure Cosmos DB, Microsoft's globally distributed, low latency database service to perform the synchronization. The service is not limited to Azure; it is applicable to all Object Storage providers supported by the Cloud File Gateway.

The following are required to add the Multi-Site Sync service:

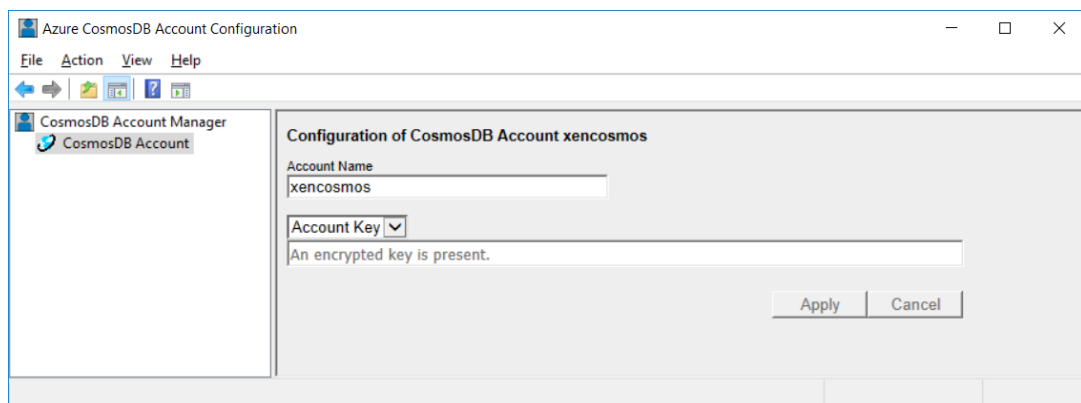
1. Purchase a Cloud File Gateway with Multi-Site Sync subscription. XenData will then provide subscription activation codes and a Cosmos DB account and key combination.
2. Install the Cloud File Gateway and Multi-Site Sync Extensions to Archive Series software.
3. License the Cloud File Gateway and Multi-Site Sync subscriptions using the License Administration Utility to enter the activation codes provided by XenData.
4. Configure cloud object storage account access for one of the cloud storage providers supported. by XenData such as AWS S3, Azure Blob Storage or Wasabi S3.

5. Add Cosmos DB Account Access, as described in the next section.

8.9.2 Adding Cosmos DB Account Access

Before starting this step it is recommended that you follow the steps outlined in [Configuring Azure Storage Accounts](#), [Configuring Amazon S3 Endpoint](#) or [Configuring a Wasabi S3 Account](#), as access to an fully configured object Storage account is required for XenData Multi-Site Sync.

1. Launch the Azure CosmosDB Account Configuration utility as follows:
 1. Click the Windows Start icon.
 2. Open the XenData program group
 3. Click the Azure **CosmosDB Account Configuration** entry in the list.
2. Left-click on 'CosmosDB Account'.



3. Enter the name of the CosmosDB Account.
4. Enter the Account Key and click 'Apply'.
5. Restart the XenData Archive Series service

Once you have completed the recommended actions outlined at the top of this page, along with the instructions here, on each of your XenData servers, be they physical or virtual, you will be ready to use the XenData Multi-Site Sync.

8.10 S3 Server Interface

8.10.1 Adding the S3 Server Interface

An overview of the S3 Server Interface is given in [Overview of S3 Server Interface](#)

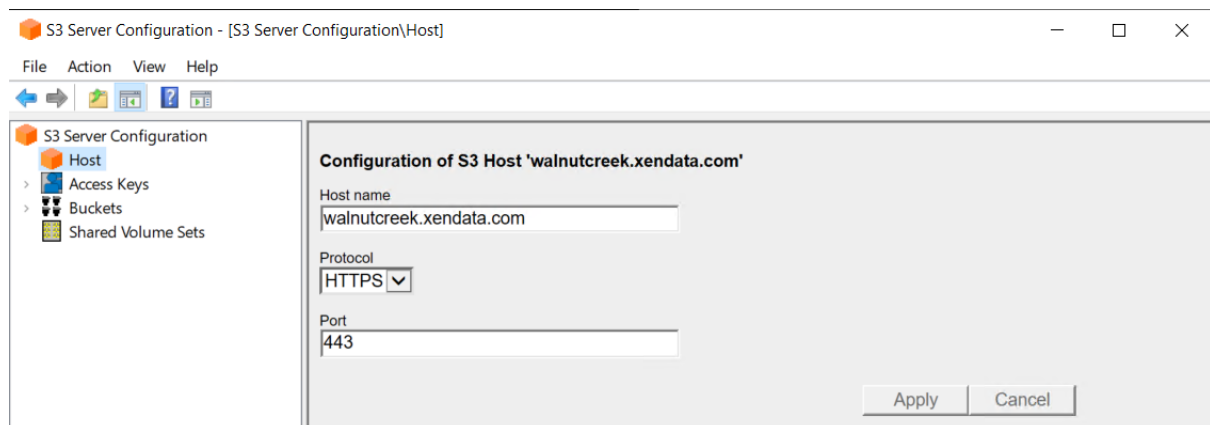
The following are required to add the S3 Server interface to an existing XenData installation.

- ❖ Purchase an XenData S3 Server Interface license, and XenData will provide an activation code.
- ❖ To be able to use HTTPS, install an SSL certificate on the Windows Server running the XenData software
- ❖ Upgrade your XenData installation with the S3 Server Interface feature, found in the installer package.
- ❖ License the S3 Server Interface, by entering the provided activation code into the XenData License Administration Utility.
- ❖ Configure the S3 Server Interface as described in the next section.

8.10.2 Configuring the S3 Server Interface

8.10.2.1 Host

1. Launch the S3 Server Configuration utility as follows:
 - a. Click the Windows Start icon.
 - b. Open the XenData program group.
 - c. Click the S3 Server Configuration entry in the list.
2. Left click on 'Host'. and you will be presented with three required variables.
 - a. Host Name: The address which external applications will use to contact the S3 server. This must be a valid DNS name.
 - b. Protocol: The protocol which will be used to communicate with the server. HTTP can be used for local connections, but HTTPS should be used for a secure connection over an external network.
 - c. Port: The port used for communication. The default HTTP port is 80, and HTTPS is 443, however these may be different depending on individual configurations.



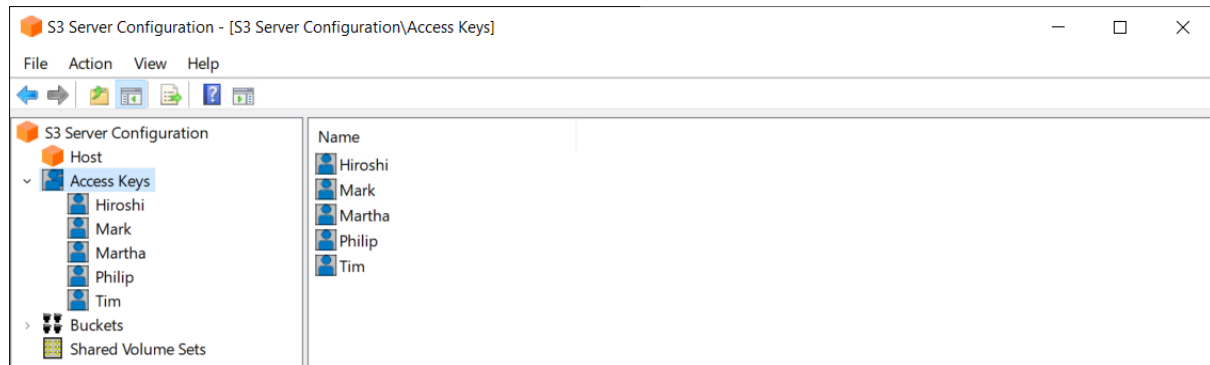
3. After the variables have been entered, left click on 'Apply',.

8.10.2.2 Access Keys

Access Keys are used to allow external applications to authenticate their access to the XenData S3 Server Interface. To generate a new access key:


1. Right click on 'Access Keys', selecting 'New' > 'Access Key'. This will create an Access Key ID, which will be created with a generated name, but it can be renamed as illustrated below.

2.



3. If required, rename the access key in the left-hand pane to create a user-friendly alias. This is performed with a right-click on the access key in the left pane.
4. Left click on the new access key in the left-hand pane, or double click on it in the right-hand pane to open it. You will see two variables here.
 - a. Access Key ID: This is the S3 access key which is the first half of the credentials used to access the S3 server from an S3 client application.
 - b. Secret Key: This is the secret key, it can only be viewed once, upon generation, but can be re-generated as many times as needed. This is the second half of the credentials used to access the S3 server from an S3 client application.

Select 'Generate' and you will see a Secret Key is generated for you. Make a copy of it, as it will not be visible again, and a new one will need to be generated if it is forgotten. This illustrated below for an Access Key that has been given the alias 'Hiroshi'.



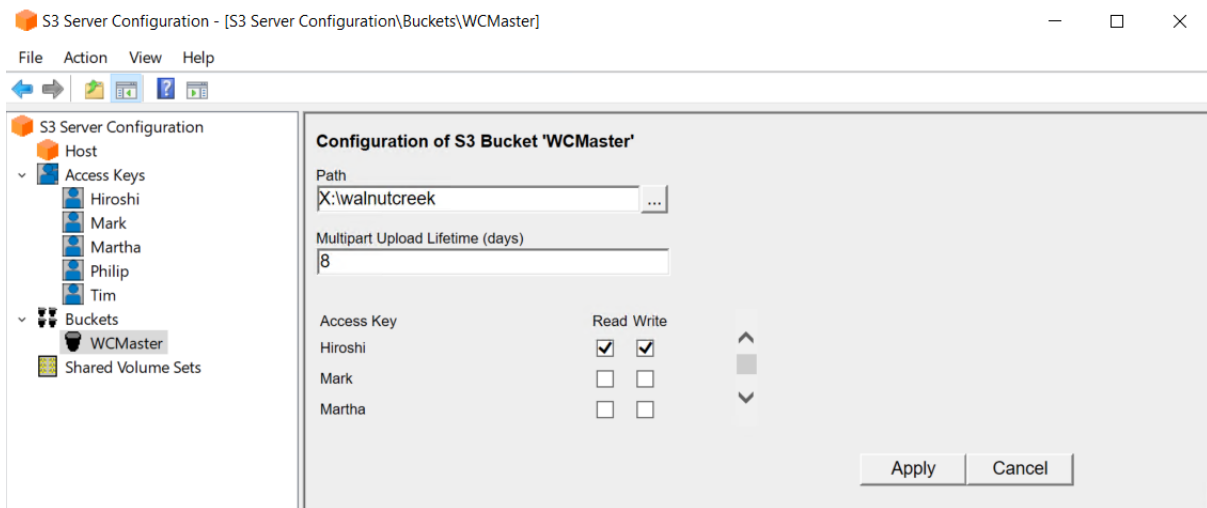
5. Once satisfied with the Access Key and Secret Key, select 'Update'.
6. Repeat the above steps to create additional access key and secret key pairs.

8.10.2.3 Buckets

Like other S3 implementations, the XenData S3 interface exposes itself to the outside world as one or more buckets. You can have as many or as few buckets as you like, and each can relate to a specific path on the XenData archive. The access to these buckets can be controlled by allowing read or write access to sets of access keys.

To configure a new bucket:

1. Right click on 'Buckets' in the left-hand pane and select 'New' > 'S3 Bucket'.
2. Give the bucket a name (an 'identity') and select 'OK' to continue.
3. Left click on 'Buckets' or expand it, to see the new bucket.
4. Left click on the new bucket in the left-hand pane, or double click on it in the right-hand pane to open it. You will see two variables and check boxes.
 - a. Path: The folder which corresponds to this bucket.
 - b. Multipart Upload Lifetime: The number of days that file upload chunks are kept in the instance that an upload does not fully complete.
 - c. Access Key Read / Write: All Access Keys are shown, with the option to enable read or write access on a per key basis. If an Access Key has been given an alias, this will be displayed.



5. Once the variables have been set, select 'Apply'.
6. Restart the XenData S3 Service, or reboot your system, and the bucket(s) will then be ready to use.

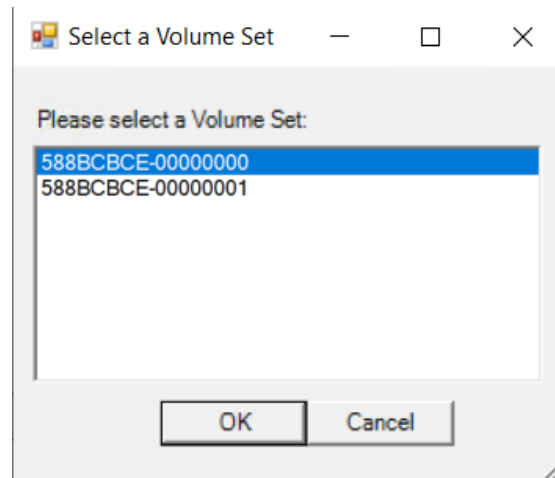
8.10.2.4 Synchronizing Multiple Archives using Shared Volume Sets

An overview of archive synchronization is given in [Synchronizing Multiple Archives](#).

To synchronize multiple XenData archives, a subscription to the XenData Global Sync service is required and the XenData S3 Interface must be installed and licensed on each archive system. The administrator for each archive system defines which of its Volume Sets will be shared.

To share out a Volume Set:

1. Right click on 'Shared Volume Sets' in the S3 Server Configuration interface.
2. Select 'New' > 'Shared Volume Set', and select the Volume Set to be shared, before clicking OK.



3. Left click on the new shared Volume Set in the left-hand pane, or double click on it in the right-hand pane to open it.
4. Select the 'Read' check box for the access keys that are to be used on the client system and click 'Apply'.



5. Restart the XenData S3 Service, or reboot your system, and the shared Volume Set(s) will then be available.

9. File Explorer Extensions

On the computer running Archive Series software, the capabilities of Windows File Explorer are extended to provide the following functionality:

- ❖ [Flushing of Files and Folders](#)
- ❖ [Pre-fetching of Files and Folders](#)
- ❖ [Smart Copy and Paste](#)
- ❖ [Enhanced Properties](#)
- ❖ [Volume View](#)
- ❖ [History Explorer](#)

For files that are stored on object storage that offers tiering, such as from AWS and Azure, the Enhanced Properties may be used to change the storage tier for a file.

9.1 Flushing of Files and Folders

Selected files and the contents of selected folders can be flushed from the disk cache using the Windows Explorer Flush option. Flushing will only occur for files that have been successfully written to LTO, ODA or object storage. The Explorer Flush option overrides the Disk Retention Rules described in [Selecting Disk Retention Rules](#).

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 disk cache, but the file is still visible and accessible to applications by restoring from LTO, ODA or object storage. The Windows offline attribute is set for all files that have been flushed.

To flush files using File Explorer:

1. Open Windows File Explorer on the computer running the Archive Series software or a connected client running the Client Utilities.
2. Select and then right-click on the required files and folders.
3. Select Flush.

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

9.2 Pre-fetching of Files and Folders

Selected files and the contents of selected folders can be pre-fetched to the cache disk cache using the Windows File Explorer Prefetch option. The Explorer Prefetch option overrides the Disk Retention Rules described in [Selecting Disk Retention Rules](#). 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 File Explorer as described in [Flushing of Files and Folders](#).

To prefetch files using Windows File Explorer:

1. Open Windows File Explorer on the computer running the Archive Series software or a connected client running the Client Utilities.
2. Select and then right-click on the required files and folders.
3. Select Prefetch.

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

Note that when using Windows File Explorer on the computer running the Archive Series software and if only a single file is selected, a Recall option is also available. This is similar to the Prefetch operation but additionally provides an on-screen display of any applicable error messages.

9.3 Smart Copy and Paste

Smart Copy and Paste is a function useful to users running the LTO Server Edition of Archive Series software. It restores files in an optimized order from LTO or ODA cartridges. It offers no significant benefit when restoring from Object Storage.

The standard copy and paste operations available within Windows File 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 functions offer two alternative methods for restoring selected files from tiered storage in an optimized order which minimizes total restore time from LTO or ODA cartridges.

To Restore Files using Smart Paste:

1. Open Windows File 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 File Explorer
2. Select, right-click and drag the selected files and folders to the required restore location.

3. Unclick and then select Smart Copy.

9.4 Enhanced Properties

9.4.1 Obtaining Enhanced Properties

Enhanced properties are available for the logical drive managed by the Archive Series software as described below.

To obtain Enhanced Properties:

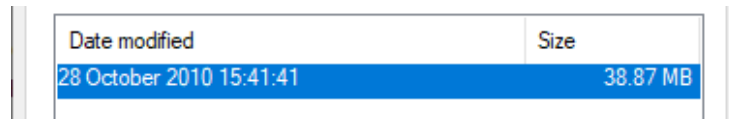
1. Open Windows File Explorer on the computer running the Archive Series software.
2. Right click on the logical drive letter under XenData control.
3. Select Properties and then select the XenData tab.

9.4.2 Changing the Object Storage Tier for a File

The XenData enhanced properties tab gives the ability to change the tier of object storage for a file. This is only applicable for files stored on object storage that supports multiple tiers, such as from AWS and Azure.

To find the storage tier options:

1. Open Windows File Explorer on the computer running the Cloud File Gateway software.
2. Right click on a file on the logical drive letter under XenData control.
3. Select Properties and then select the XenData tab.



Date modified	Size
28 October 2010 15:41:41	38.87 MB

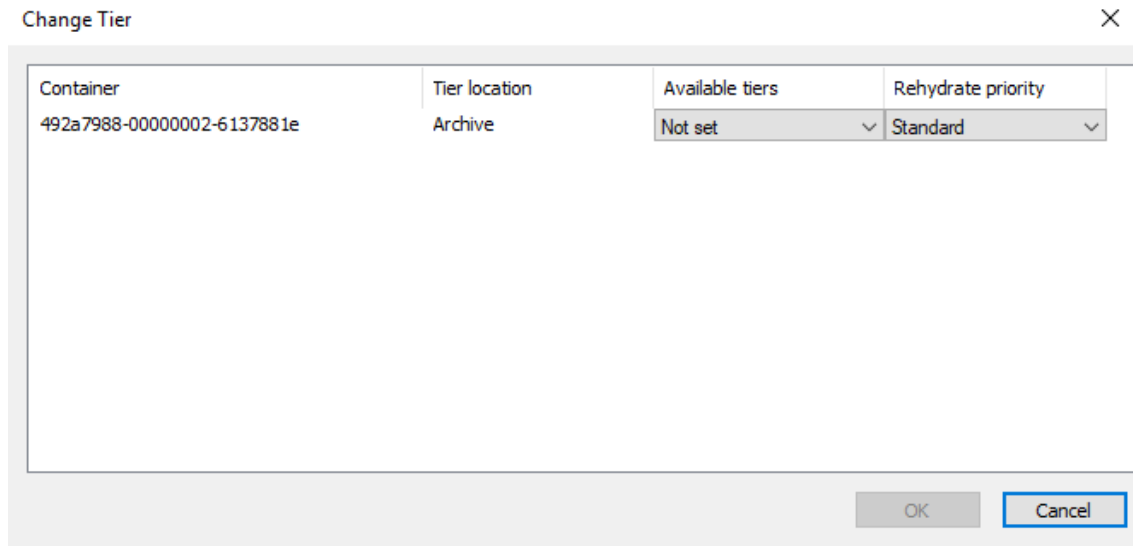
4. Select the version of the file you would like to change the tier of.

Tier location


Archive

Change Tier

5. Select 'Change Tier' to see options for moving the file to another storage tier.



- a. Container: This is the Object storage container or bucket which contains your selected file.
- b. Tier Location: The current storage tier the file is in.
- c. Available Tiers: The drop-down menu contains a list of tiers that it is possible to move the current file to. Each object storage service has different storage tiers available, with their own costs.



- d. Rehydrate Priority: If restoring a file from the Archive Tier, Glacier or Deep Glacier, the Rehydrate Priority drop-down will appear. This allows the selection of the rehydration priority. Each object storage service has different rehydration priorities available, with their own costs and time scales.
6. After the selection of a storage tier and rehydration policy (if rehydrating). Pressing 'OK' will begin the process immediately. The tier of a file currently undergoing rehydration will specify that it is undergoing rehydration, and the tier to which it is rehydrating.

9.5 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 File Explorer:

1. Open Windows File Explorer on the computer running the Archive Series software.
2. Select Volume View in the left navigational pane.
3. Browse the Volume View.

9.6 History Explorer

History Explorer is used with the LTO Server Edition of Archive Series software and the Optical Disc Archive Extension to obtain the complete history and status of any file that the system knows about. It 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.

The default behavior of the Cloud File Gateway Extension does not retain old file versions and deleted files on Object Storage. However, it may be configured to retain old versions of files and deleted files via a registry setting. In this case, History Explorer may be used to obtain the complete history and status of any file written to Blob Storage, as it does for files written to LTO or ODA cartridges.

To browse with History Explorer

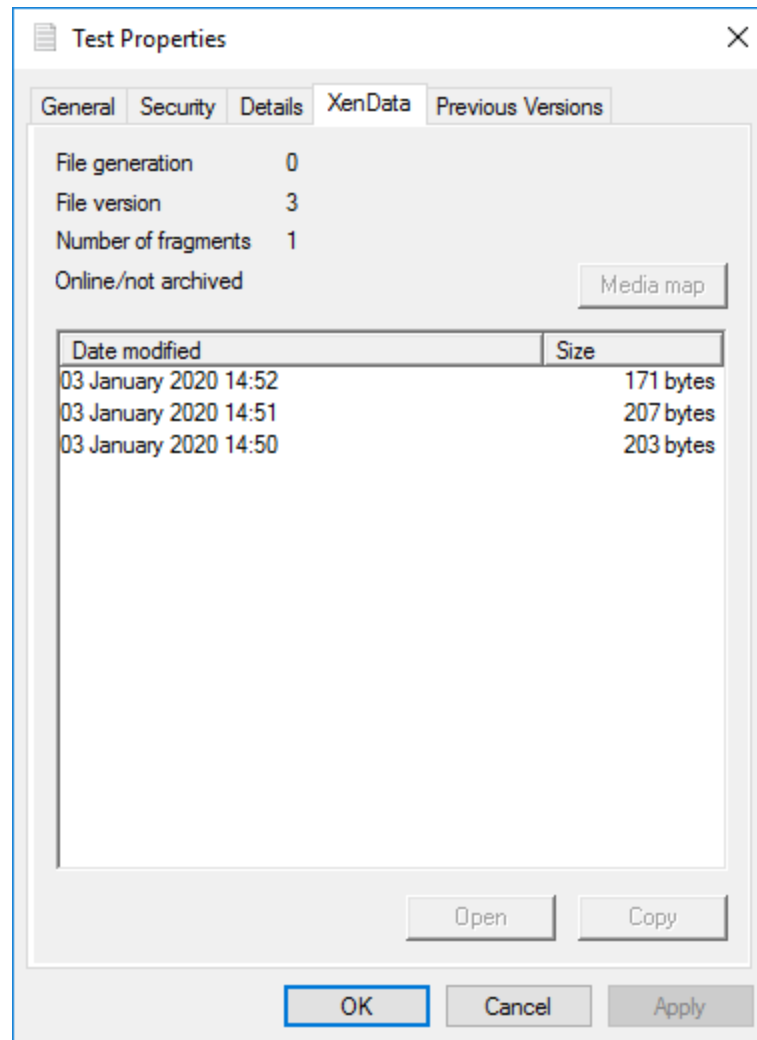
1. Open Windows File Explorer.
2. Select **History Explorer** in the left navigational pane.
3. Browse the file system.

History Explorer allows you to browse the logical drive managed by the XenData software and will show current (i.e. non-deleted) and deleted files. Any deleted files are shown as greyed out.

For the logical drive managed by XenData, Archive Series software adds a XenData tab to a file's properties dialog when browsing using either the standard Windows File Explorer view or using History Explorer. The XenData tab will identify all file versions and by clicking a version to highlight it and then clicking Open or Copy, that file version may be opened or copied to another storage location.

To view the versions of a file for either current or deleted files:

1. Open Windows File Explorer.
2. Select **History Explorer** in the left navigational pane.
3. Browse the file system.
4. Select and then right-click on the required file.
5. Select **Properties**.
6. Select the XenData tab.



To open an old version of a 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 Archive Series Windows Explorer extension.)

To restore an old version of a 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.

10. Metadata Backup

The Metadata Backup program backs up and restores:

- ❖ File system metadata which is stored on the cache disk
- ❖ The State File which contains Volume information and the Tiered Storage Management Console settings, including File Group and Volume Set configuration settings.

Metadata backups may be configured as scheduled tasks as described in the [Scheduler section](#).

10.1 About Metadata Backup

If a XenData Archive Series system has to be rebuilt, perhaps due to failure of the cache disk, the file system metadata may be rebuilt by using the **Build Catalog**, **Import Folder Structure** and **Import Data** functions available in the Tiered Storage 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 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.

Metadata backups can be scheduled using the XenData Scheduler, as described in [Scheduling Metadata Backup](#).

10.2 Starting Metadata Backup

To start the program:


1. Click the Windows Start icon
2. Open the XenData program group
3. Click the **XenData Metadata Backup** entry in the list

10.3 Selecting Backup or Restore

The Metadata Backup program performs two types of operation:

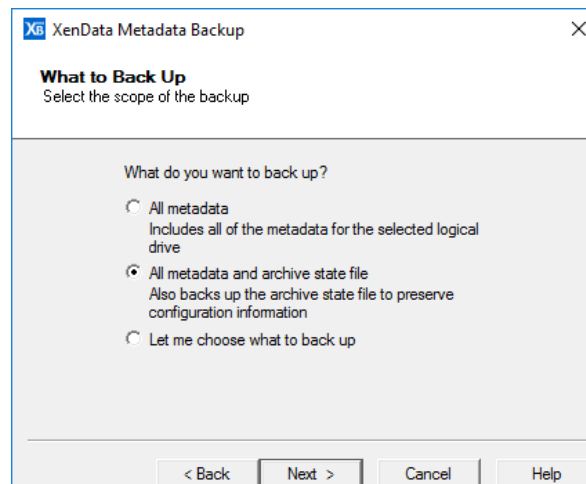
- ❖ Make backup - makes a backup of the metadata in the system in its current state. See 'Making a Predefined backup' or 'Making a Custom Backup' below.
- ❖ Restore from backup – restores metadata from a backup file onto the cache disk volume. See 'Restore from backup' below.

Select the desired option and click Next to continue.



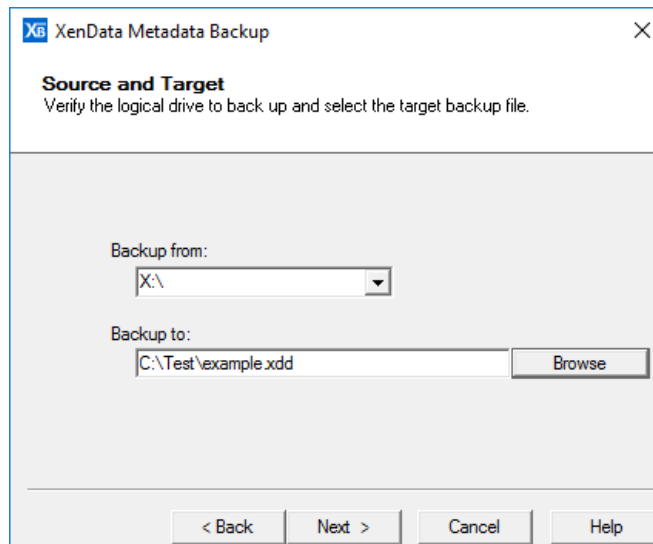
10.4 Making a Predefined Backup

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**.

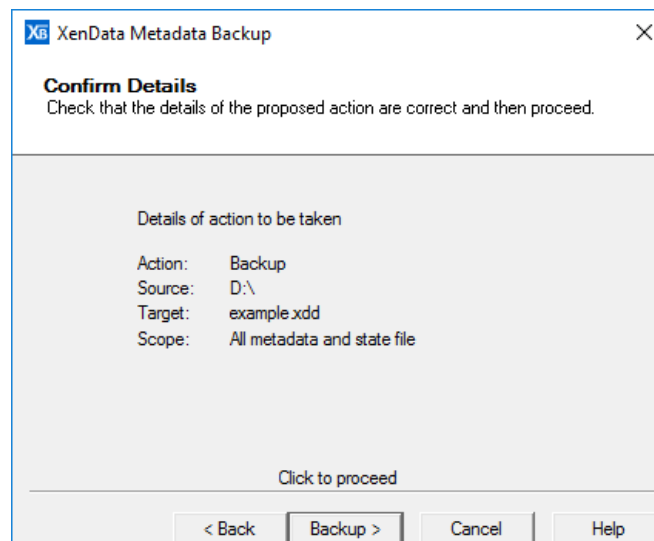


There are two predefined backup types. **All metadata** will back up all the file system metadata, and **All metadata and XenData state file** will also include the XenData state file.

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




1. Verify that the logical drive letter to be backed up is correct.
2. Specify the output path and file name. The output file name should be inserted in the **Backup to** edit box. Click **Browse** to assist in specifying the path and file name.
3. Click **Next** to continue.

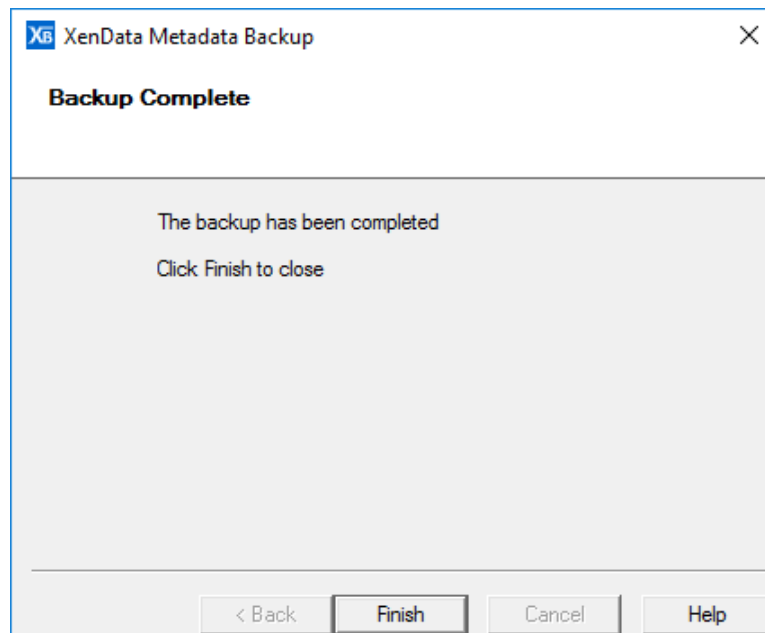


The next page presents the details of the backup, and gives the option to go back and correct if necessary.

1. Verify the backup details.
2. Click **Backup** to perform the backup.
3. A progress dialog box appears that shows the backup progress, as illustrated below.



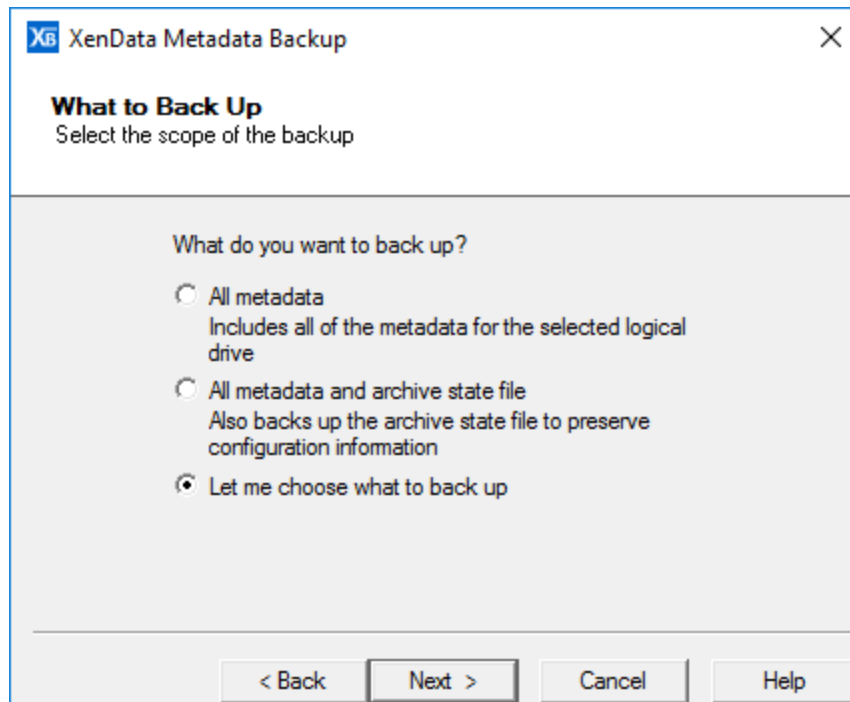
If the backup completed successfully, you will be presented with a confirmation page saying Backup Complete. Click **Finish** to dismiss the dialog and exit the program.



10.5 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 XenData 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 logical 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.



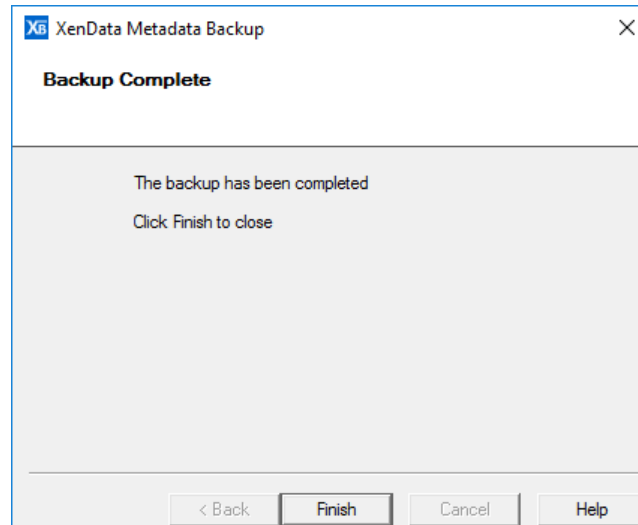
This page presents the details of the backup, and gives the option to go back and correct if necessary.

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.



10.6 Restoring 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 XenData state file.

Either start the Metadata Backup program, select **Restore from backup** and click **Next** on the starting page, or double click on a 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).

To restore a metadata backup:


1. Specify the input backup file to restore from, or verify that the correct file name has been determined automatically.
2. Verify the logical drive letter 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.

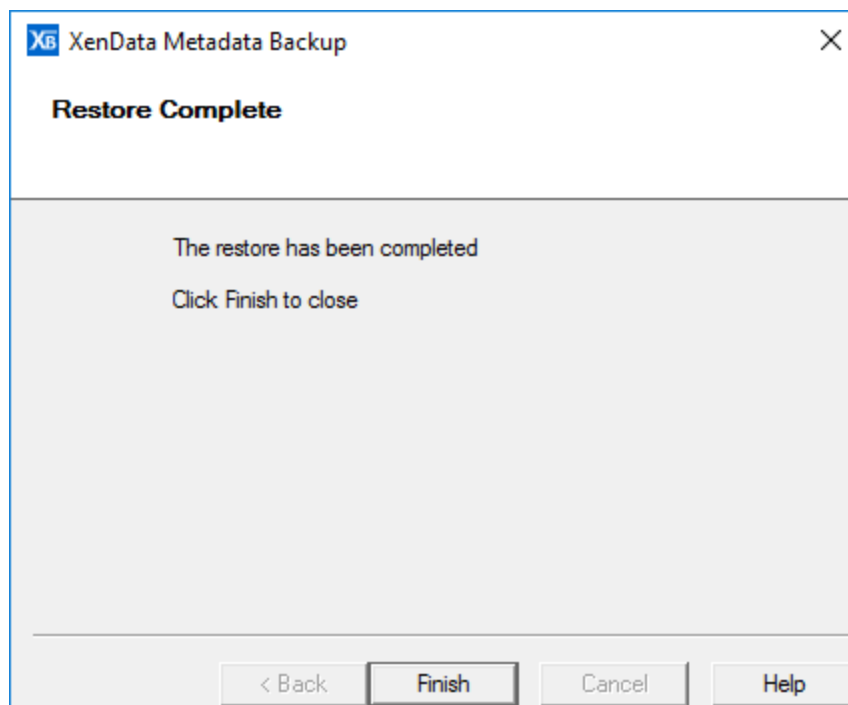
A progress dialog box will appear so that you can check the status of the restore operation. If the option to **Ask before overwrite** was selected during restore configuration, dialog boxes similar to the one shown below might appear, asking if existing metadata should be overwritten, and giving a category 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 completes successfully, you will be presented with a confirmation page saying Restore Complete. Click **Finish** to dismiss the dialog box and exit the program.



11. 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 XenData state file. It does not support scheduling of custom backups.
- ❖ Deferred Writing which defers the initial writing of files to a Volume and allows you to specify a scheduled time period when data can be written to LTO, ODA or object storage. It is useful for prioritizing file restore operations during times of peak demand.
- ❖ Replication which allows scheduling of replicated Volume Sets.
- ❖ FS Mirror which is an upgrade option that is licensed separately. It provides replication and synchronization of file systems accessible to the server running Archive Series Software.
- ❖ FS Mirror Test is an upgrade option included in the FS Mirror license. It provides a way of testing the replication and synchronization of file systems accessible to the server running Archive Series Software.

11.1 Starting the Scheduler

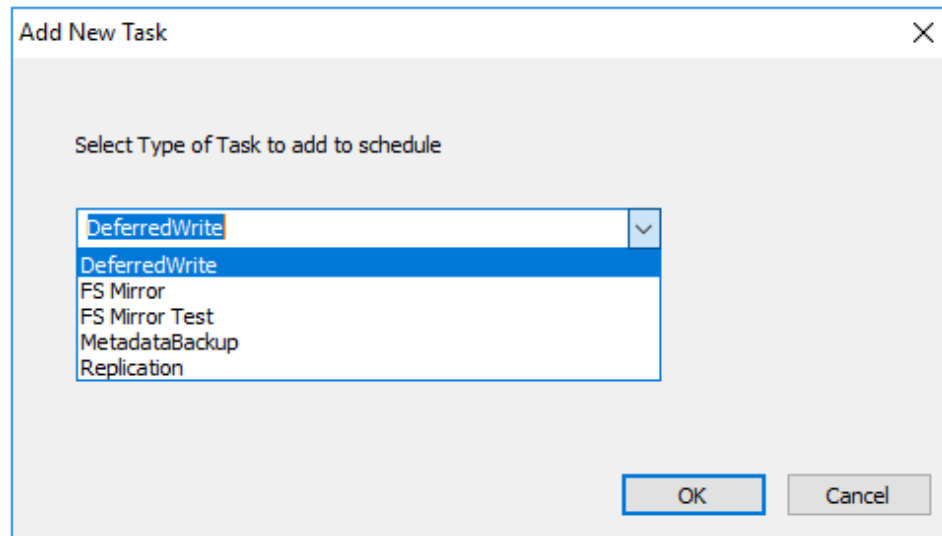
To start the Scheduler:

1. Click the Windows Start icon.
2. Open the XenData program group
3. Click the **XenData Scheduler** entry in the list

11.2 Adding a Task

To Add a Task:

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



11.3 The Scheduler Status Display

An example of the Scheduler status display is shown below.

Task Type	Task Name	Status	Last Run Time	Last Run Status	Next Run Time	Expires	Recurrence
MetadataBackup	Weekly Backup	Idle	2020-01-03 15:20	OK	2020-01-10 15:20	--	Weekly
DeferredWrite	After Hours Up...	Idle	--	--	2020-01-03 21:15	--	Daily
FS Mirror	FS Mirror	Running	2020-01-03 15:35	--	2020-01-03 15:45	--	Daily

The display columns are as follows:

- ❖ Task Type - standard options are Metadata Backup, Deferred Write and Replication. FS Mirror and FS Mirror Test will appear as an option if the FS Mirror software has been installed and its license activated.
- ❖ 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 – When the task is running, a green progress bar is displayed and an administrator can stop the task using the stop button.

- 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 - only runs using the Run Now option.
 - Hourly - Task is run once per a specified number of hours. 1, 2, 3, 4, 8 and 12 hour options are selectable from the drop-down list.
 - 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.

11.4 Editing and Deleting Tasks

To Edit a Task

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

To Delete a Task

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

11.5 Starting and Stopping Tasks

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

To Run a Task "Now"

1. Start the 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 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.

11.6 Scheduling Metadata Backup

The screenshot shows the 'MetadataBackup' dialog box with the following configuration:

- Recurrence:** Weekly (selected)
- Start:** 2020-01-03 15:20 (Start date and time), 2022-12-31 15:20 (End date and time), Expire checkbox checked.
- Task Name:** Weekly Backup
- Stop task if it runs longer than:** Checked.
- Directory path:** C:\MetadataBackups
- Delete previous backups:** Unchecked.

Options for the Metadata Backup task are as follows:

- ❖ Recurrence is one of:
 - None -only runs using the Run Now option.
 - Hourly - Task is run once per a specified number of hours. 1, 2, 3, 4, 8 and 12 hour options are selectable from the drop-down list.
 - 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 XenData 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.

11.7 Scheduling Deferred Write

DeferredWrite

Recurrence: None Daily Weekly Monthly

Hourly:

Start: 2020-01-03 21:15

Expire: Expire 2021-12-31 21:15

Task Name: After Hours Update

Stop task if it runs longer than

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"/> 492A7988-00000000	LTO

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 - only runs using the Run Now option.
 - Hourly - Task is run once per a specified number of hours. 1, 2, 3, 4, 8 and 12 hour options are selectable from the drop-down list.
 - 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.

Options for the Media Replication task are as follows:

- ❖ Recurrence is one of
 - None - only runs using the Run Now option.
 - Hourly - Task is run once per a specified number of hours. 1, 2, 3, 4, 8 and 12 hour options are selectable from the drop-down list.
 - 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.

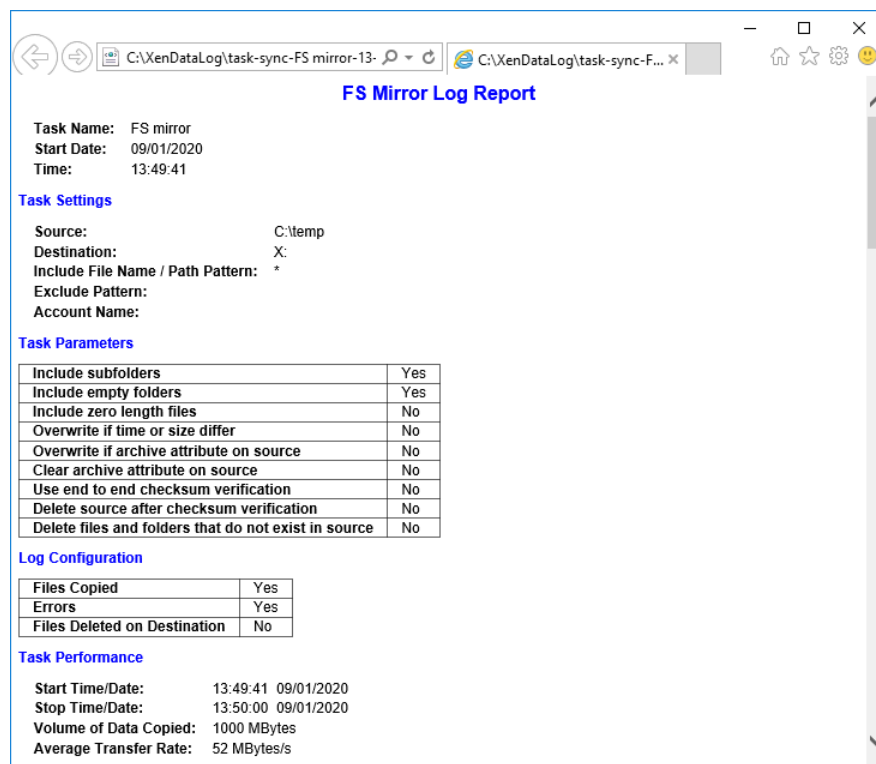
- ❖ Volume Sets with Replication Enabled - is a list of all the Volume Sets in the system that have replication enabled. The Volume Sets that are selected with a check mark are controlled by this particular replication task.

11.9 Scheduling FS Mirror

Options for the FS Mirror task are as follows:

- ❖ Recurrence is one of
 - None – only runs using the Run Now option.
 - Hourly - Task is run once per a specified number of hours. 1, 2, 3, 4, 8 and 12 hour options are selectable from the drop-down list.
 - 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.
- ❖ Use Log File – checking this box creates a log file for the task each time it runs. Logs are found in %SystemRoot%\XenDataLog\, typically C:\XenDataLog\, and are in an XML format. When you double-click on the xml file, it launches a human-readable html report.
 - Log all copied files – records all files copied in the log file. If more than 10,000 files are expected to be copied, it is advisable not to enable this option, as it will make opening the log file very slow.
 - Log Errors – records, in the log file, all files not copied due to an error. If a file was open at the time that the task ran, it would not be copied and would be recorded as an error.
 - Log files deleted on destination – records, in the log file, all files deleted on the destination.
 - Log files deleted on source – records, in the log file, all files deleted on the source.



- ❖ Source Folder – the folder which contains the source files and folders.
- ❖ Destination Folder – the folder where the files and folders will be copied.

- ❖ Include file name or file path pattern – a required parameter which controls which files will be copied based on a pattern match. The default value is '*', which copies all files, as long as they match the check box settings.
- ❖ Exclude Pattern – an optional parameter that determines files to be excluded from the copy, regardless of other rules, like the previous setting, it is based on a pattern match. An example would be '.tmp', which would exclude all files with the .tmp extension. The separator between multiple exclusions is ;. An example with multiple exclusions would be '.tmp;.mxf;.png', which would exclude all files with a .tmp, .mxf or .png extension. Spaces should not be included.
- ❖ User Account – an optional parameter, only needed if you are copying across a network that requires user authentication. Takes standard domain\user account credentials.
- ❖ Password – an optional parameter, only needed if you are copying across a network that requires user authentication. This is the password for the User Account.
- ❖ Include subfolders – checking this box will make the FS Mirror task include source subfolders and their contents.
- ❖ Include empty folders – checking this box will make the FS Mirror task include folders which contain no files.
- ❖ Include zero length files – checking this box will make the FS Mirror task include files which contain no data, and as such have no size.
- ❖ Overwrite if size or time differ – checking this box will make the FS Mirror task overwrite files at the destination if they have the same name, but a different size or modification time to those in the source.
- ❖ Overwrite if source has archive attribute set – checking this box will make the FS Mirror task overwrite files at the destination, if the source file has the archive attribute set.
- ❖ Clear archive attribute on source – checking this box will make the FS Mirror task remove the archive attribute from the source file after it has been copied successfully.
- ❖ Use end-to-end checksum verification – this option can only be enabled when the destination is a XenData LTO or object storage archive. Checking this box will make the FS Mirror task utilize end-to-end checksum verification. Before and after each file is copied, a checksum will be performed. This ensures that the file that reaches the destination is the same as that which leaves the source. To enable, you will need to have Logical Block Protection enabled within the XenData Tiered Storage Management Console.
- ❖ Delete source after checksum verification – this option can only be enabled when 'Use end-to-end checksum verification' is enabled. With this option enabled, the source files will be deleted after the checksum verification has confirmed that the file has been completely written to the destination with byte-for-byte accuracy.

- ❖ Delete files and folders that do not exist in source – checking this box will make the FS Mirror task delete all files and folders at the destination that do not exist in the source folder.
- ❖ Check - launches a test of the current task, to determine the result of the current settings. It does not make any changes to files and folders in either the source or destination folders. The test run will inform the user of any files which would not be copied, along with a reason, which can be useful for modifying the task before running. Unlike FS Mirror Test, this test does not use the authentication parameters defined by the User Account and Password settings; it uses the authentication applicable to the currently logged in user.

11.10 Scheduling FS Mirror Test


The screenshot shows the 'FS Mirror Test' configuration window. It includes the following fields and options:

- Recurrence:** Radio buttons for None (selected), Hourly (1 hour), Daily, Weekly, and Monthly.
- Start:** Two date and time pickers, both set to 2020-01-06 08:57.
- Task Name:** Text field containing 'FS Mirror Test'.
- Log and Error Options:**
 - Use Log File
 - Log Errors
 - Log all copied files
 - Log deletions on destination
- File Handling Options:**
 - Include subfolders
 - Include empty folders
 - Include zero length files
 - Overwrite if size or time differ
 - Overwrite if source has archive attribute set
 - Delete files and folders that do not exist in source
- Folder Paths:**
 - Source Folder: C:\Test
 - Destination Folder: X:\Test
- Patterns:**
 - Include file name or file path pattern: *
 - Exclude Pattern: (empty)
- Authentication:**
 - User Account: (empty)
 - Password: (empty)
- Buttons:** Save, Cancel

The FS Mirror Test task does not make any changes to files and folders in either the source or destination folders. It identifies what files and folders would be copied and/or deleted if these settings were made using an FS Mirror run. It uses the authentication parameters defined by the User Account and Password settings for the task.

Options for the FS Mirror Test task are as follows:

- ❖ Recurrence is one of
 - None – only runs using the Run Now option.
 - Hourly - Task is run once per a specified number of hours. 1, 2, 3, 4, 8 and 12 hour options are selectable from the drop down list.
 - 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.
- ❖ Use Log File – Use Log File – checking this box creates a log file for the task each time it runs. Logs are found in %SystemRoot%\XenDataLog\, typically C:\XenDataLog\, and are in an XML format. When you double-click on the xml file, it launches a human-readable html report.
 - Log all copied files – records, in the log file, all files that would be copied. If more than 10,000 files are expected to be copied, it is advisable not to enable this option, as it will make opening the log file very slow.
 - Log Errors – records, in the log file, all files that would not copied due to an error. If a file was open at the time that the task ran, it would not be copied and would be recorded as an error.
 - Log files deleted on destination – records, in the log file, all files that would be deleted on the destination.



- ❖ Source Folder – the folder which contains the source files and folders.
- ❖ Destination Folder – the folder where the source files and folders will be copied.
- ❖ Include file name of file path pattern – a required parameter which controls which files would be copied based on a pattern match. The default value is '*', which copies all files, as long as they match check box settings.
- ❖ Exclude Pattern – an optional parameter that determines files to be excluded from the copy, regardless of other rules, like the previous setting it is based on a pattern match. An example would be '.tmp', which would exclude all files with the .tmp extension.
- ❖ User Account – an optional parameter, only needed if you are copying across a network that requires user authentication. Takes standard domain\user account credentials.
- ❖ Password – an optional parameter, only needed if you are copying across a network that requires user authentication. The password for the previously mentioned user account.
- ❖ Include subfolders – checking this box will make the task include source sub-folders and their contents..
- ❖ Include empty folders – checking this box will make the task include folders which contain no files.

- ❖ Include zero length files – checking this box will make the task include files which contain no data, and as such have no size.
- ❖ Overwrite if size or time differ – checking this box will make the task include files to be overwritten at the destination if they have the same name, but a different size or modification time to those in the source.
- ❖ Overwrite if source has archive attribute set – checking this box will make the task overwrite files to be overwritten at the destination at the destination, if the source file has the archive attribute set.

12. Reports

The Report Generator allows you to create, save and restore a range of different reports about the files managed by the system.

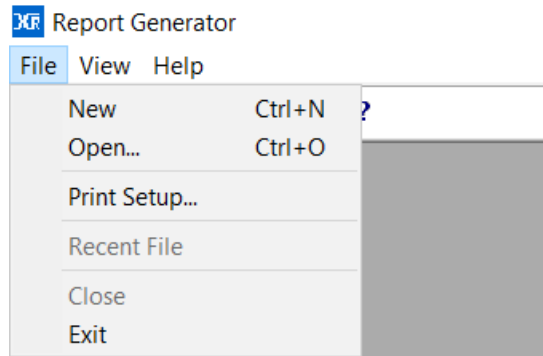
12.1 Starting the Report Generator

1. Click the Windows Start icon.
2. Open the XenData program group
3. Click the **XenData Report Generator** entry in the list.


12.2 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, then select the required XRG file and open it.

12.3 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 Volume 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.

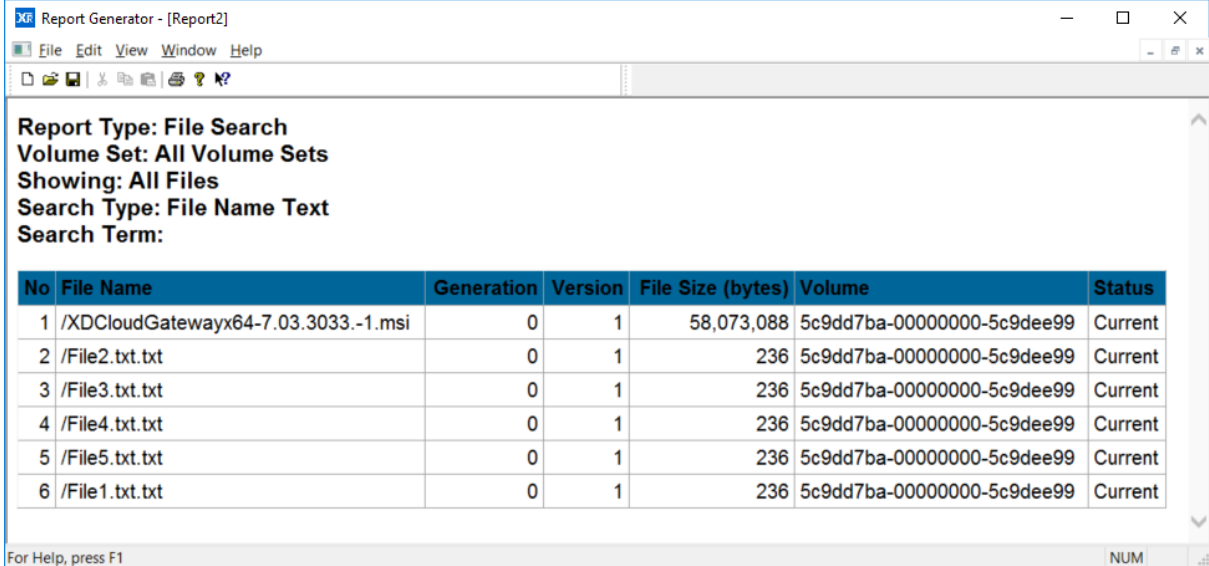
It is possible to search using a simple text search or using a Regular Expression. When **File Name Text Search** is chosen, the search option supports wild cards.

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.

12.3.1 Interpreting a File Search Report

An example of a File Search Report is shown below.



Report Generator - [Report2]

File Edit View Window Help

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

No	File Name	Generation	Version	File Size (bytes)	Volume	Status
1	/XDCloudGatewayx64-7.03.3033.-1.msi	0	1	58,073,088	5c9dd7ba-00000000-5c9dee99	Current
2	/File2.txt.txt	0	1	236	5c9dd7ba-00000000-5c9dee99	Current
3	/File3.txt.txt	0	1	236	5c9dd7ba-00000000-5c9dee99	Current
4	/File4.txt.txt	0	1	236	5c9dd7ba-00000000-5c9dee99	Current
5	/File5.txt.txt	0	1	236	5c9dd7ba-00000000-5c9dee99	Current
6	/File1.txt.txt	0	1	236	5c9dd7ba-00000000-5c9dee99	Current

For Help, press F1 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, XenData Archive Series 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 Volume followed by the total size of the file in bytes.

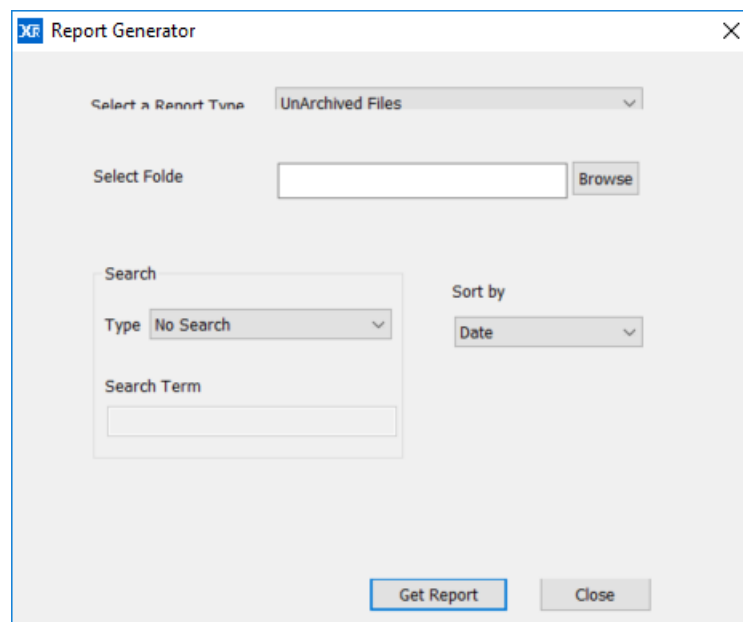
- Volume - this is the Volume that contains the file.
- Type - The status of the file is displayed as one of the following:
 - Current - this is the most recent version of the file, accessible through the archive drive letter.
 - Renamed - the file has been renamed and is now accessible under a different 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**.
 - Rearchived - the file has been rearchived/repacked, and the target volume is displayed

12.4 UnArchived Files Report

The UnArchived Files Report lists files which are not fully archived to 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**.

12.4.1 Interpreting an UnArchived Files Report

An example of an UnArchived Files Report is shown below.

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

No	File Name	Generation	Version	Replica	Volume	Status
1	/AFile1.txt.txt	0	1	1	Unknown	Not Archived
2	/AFile2.txt.txt	0	1	1	Unknown	Not Archived
3	/AFile3.txt.txt	0	1	1	Unknown	Not Archived
4	/AFile4.txt.txt	0	1	1	Unknown	Not Archived
5	/AFile5.txt.txt	0	1	1	Unknown	Not Archived

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, XenData Archive Series 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 - when a file is written to a replicated volume set, a copy of that file will be written to each of the tapes in the replicated volume set. This column tells you if there is a replica, and how many replicas there are.
- Volume - available in cases where a Volume 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 it is not archived properly. The status of the file instance is displayed as one of the following:
 - Not Archived - the file is not archived in a Volume
 - Partially Archived - the file is not fully archived.
 - Unverified Archived - the file data was written to a Volume, but Cloud File Gateway software was unable to verify that the operation had completed successfully.
 - Archived - this instance of the file is archived correctly.

12.5 Volume Contents Report

The Volume Contents Report lists the contents of the Volume.

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 in the Volume 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 Volume 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.

12.5.1 Interpreting a Volume Contents Report

An example of a Volume Contents Report is shown below.

Report Generator - [Report4]

File Edit View Window Help

Report Type: Volume Contents
 Volume: 5C9DD7BA-00000000-5C9DEE99
 Showing: All Files
 Search Type: None
 Sorted by: Date

No	File Name	Generation	Version	File Size (bytes)	Date Archived	Type
1	/XDCloudGatewayx64-7.03.3033.-1.msi	0	1	58,073,088	Mar 29 2019 10:08	Current
2	/File2.txt.txt	0	1	236	Mar 29 2019 10:14	Current
3	/File3.txt.txt	0	1	236	Mar 29 2019 10:14	Current
4	/File4.txt.txt	0	1	236	Mar 29 2019 10:14	Current
5	/File5.txt.txt	0	1	236	Mar 29 2019 10:14	Current
6	/File1.txt.txt	0	1	236	Mar 29 2019 10:14	Current

Done NUM

The display columns are described below.

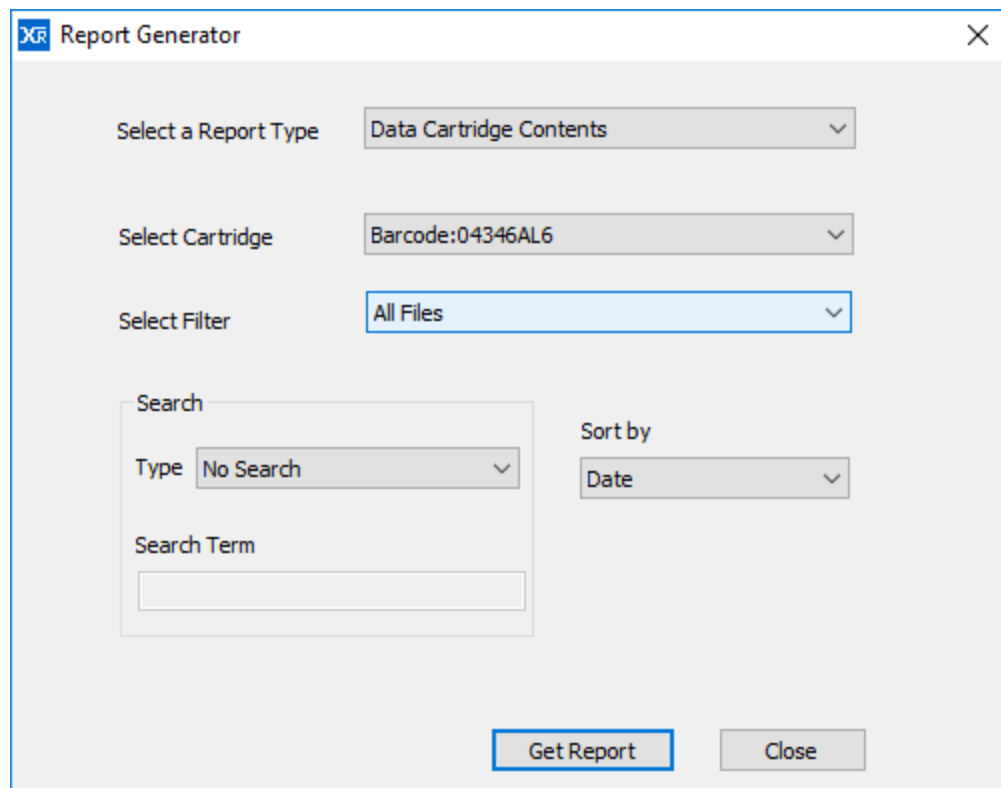
- **No** - the sequence number of the file in the display sorted by either date or file name, as defined by the **Sort by** selection.
- **File Name** - the file name including full path from the root of the archive logical drive letter.
- **Generation** - when a file of a given name and path is first 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, Cloud File Gateway 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**.
- Rearchived - the file has been rearchived/repacked, and the target volume is displayed

12.6 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.



The screenshot shows the 'XR Report Generator' dialog box. It has a title bar with 'XR Report Generator' and a close button. The main area contains three dropdown menus: 'Select a Report Type' (set to 'Data Cartridge Contents'), 'Select Cartridge' (set to 'Barcode:04346AL6'), and 'Select Filter' (set to 'All Files'). Below these are search options: 'Search Type' (set to 'No Search'), 'Search Term' (empty text box), and 'Sort by' (set to 'Date'). At the bottom are 'Get Report' and 'Close' buttons.

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.

12.6.1 Interpreting a Cartridge Contents Report

An example of a Cartridge Contents Report is shown below.

Report Generator - [Report1]

File Edit View Window Help

Report Type: Data Cartridge Contents
 Cartridge: Barcode: 04346AL6
 Showing: All Files
 Search Type: None
 Sorted by: Date

No	File Name	Generation	Version	File Size (bytes)	Date Archived	Type
1	/Test/001/001.dat	0	1	10,485,760	Jan 03 2020 16:16	Current
2	/Test/001/002.dat	0	1	10,485,760	Jan 03 2020 16:16	Current
3	/Test/001/003.dat	0	1	10,485,760	Jan 03 2020 16:16	Current
4	/Test/001/004.dat	0	1	10,485,760	Jan 03 2020 16:16	Current
5	/Test/001/005.dat	0	1	10,485,760	Jan 03 2020 16:16	Current
6	/Test/001/006.dat	0	1	10,485,760	Jan 03 2020 16:16	Current
7	/Test/001/007.dat	0	1	10,485,760	Jan 03 2020 16:16	Current
8	/Test/001/008.dat	0	1	10,485,760	Jan 03 2020 16:16	Current
9	/Test/001/009.dat	0	1	10,485,760	Jan 03 2020 16:16	Current

For Help, press F1

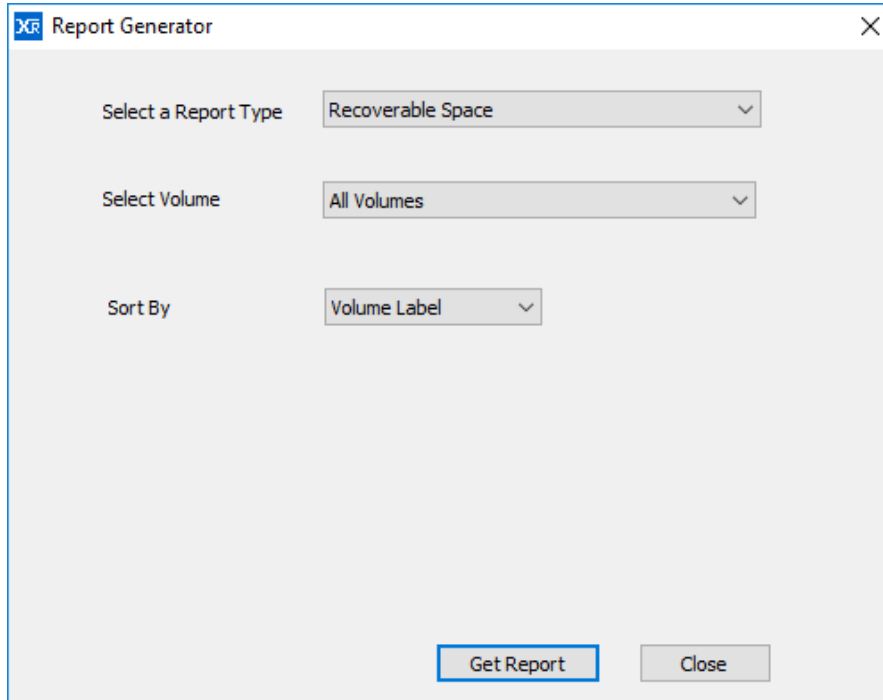
The display columns are described below.

- 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, Archive Series 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**.

12.7 Recoverable Space Report

To Run a Recoverable Space Report

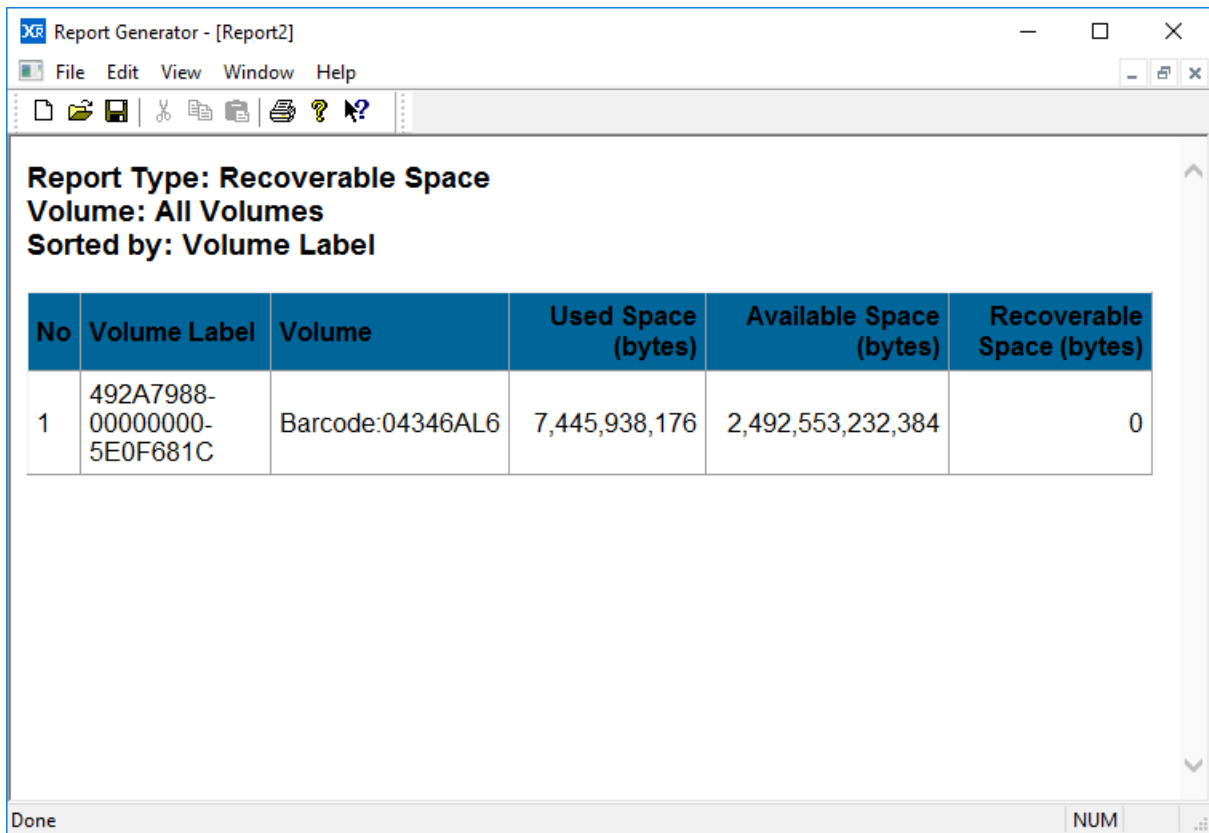
1. Start the Report Generator.
2. Select the **File** and **New** menu options.
3. Select **Recoverable Space** as the report type.



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

12.7.1 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.



Report Generator - [Report2]

File Edit View Window Help

Report Type: Recoverable Space
Volume: All Volumes
Sorted by: Volume Label

No	Volume Label	Volume	Used Space (bytes)	Available Space (bytes)	Recoverable Space (bytes)
1	492A7988-00000000-5E0F681C	Barcode:04346AL6	7,445,938,176	2,492,553,232,384	0

Done NUM

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.

13. Alert Module

The XenData Alert Module is designed for use with the Archive Series software and provides e-mail and onscreen alerts. The alerts are derived by filtering and categorizing events recorded by the Archive Series software in the Windows Event Log.

13.1 About the Alert Module

The XenData Alert Module is designed for use with the Archive Series software and provides e-mail and onscreen alerts. The alerts are derived by filtering and categorizing events recorded by the Archive Series software in the Windows Event Log.

The XenData Alert Module has two major components:

- ❖ Event Monitor with integrated e-mail notification that runs on the same computer as the Archive Series 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 Windows Event Log. For more information, see [About the Event Monitor](#).
- ❖ On-Screen Messaging is a program that runs on the same computer as the Archive Series software and may also be run on one or more Windows clients. It is installed on a Windows client using the XenData Client Utilities. The On-Screen Messaging program can be configured to display via message boxes and system tray notification. For more information see [About On-Screen Messaging](#).

13.2 About the Event Monitor

The Event Monitor runs on the same server as the Archive Series 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 Email Server](#).


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

13.3 Configuring the Event Monitor

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 Event Monitor Configuration:

1. Click the Windows Start icon.
2. Open the XenData program group.
3. Click the **XenData Event Monitor Configuration** entry in the list.



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 Email 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 message for all connected computers that are running the On-Screen Messaging program.

13.4 About Event Categories

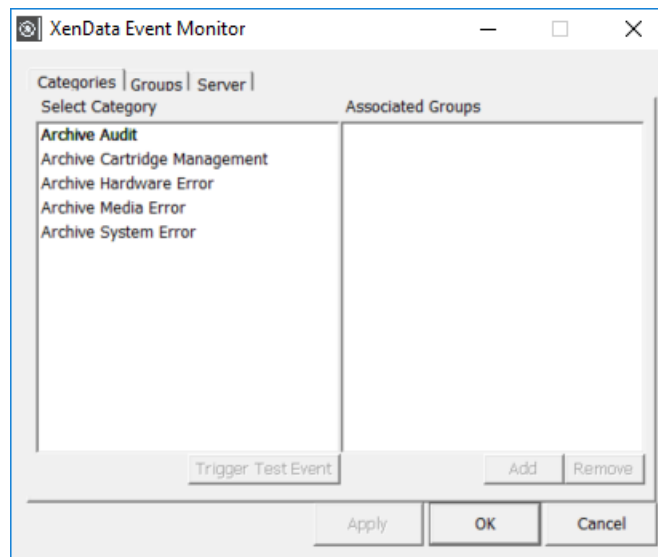
The Event Monitor is pre-configured with five Event Categories:

- ❖ **Archive Audit:** This category of event messages describes the successful completion of routine operations.
- ❖ **Archive Media Management:** This category of event messages may require routine action from the gateway operator.
- ❖ **Archive Media Error:** This event category consists of error messages associated with Volumes.
- ❖ **Archive Hardware Error:** This event category consists of error messages associated with the LTO, ODA or object storage.
- ❖ **Archive 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](#).

13.5 Configuring Event Categories

Launch the Event Monitor configuration screen by starting the configuration program as described in [Configuring the Event Monitor](#). 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**.

13.6 About Recipient Groups

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

13.7 Configuring Recipient Groups

Launch the Event Monitor configuration screen by starting the configuration program as described in [Configuring the Event Monitor](#). 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**

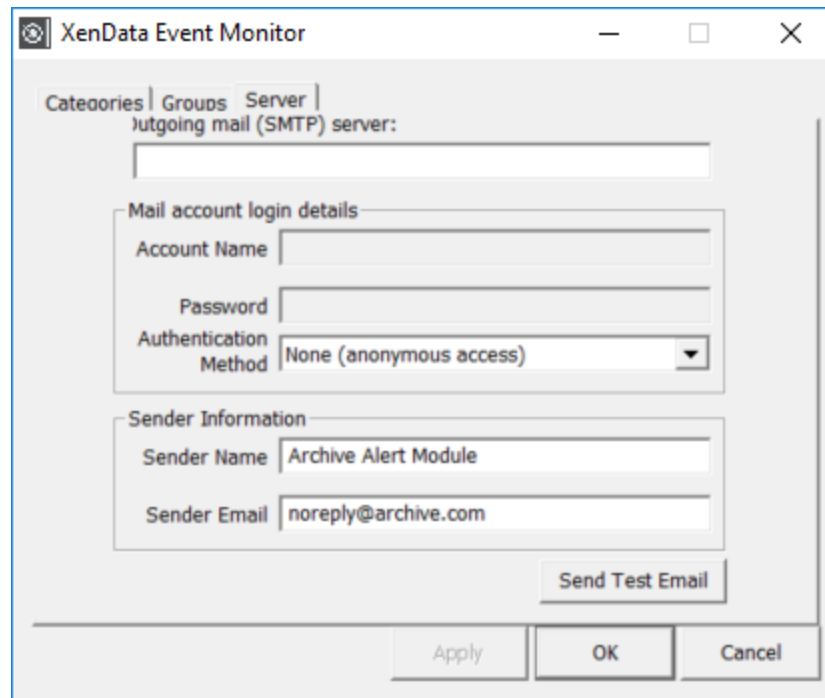
13.8 About the Email Server

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 Email server and the Email account information is described in [Configuring the Email Server](#).

13.9 Configuring the Email Server

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



The screenshot shows the 'XenData Event Monitor' window with the 'Server' tab selected. The 'Outgoing mail (SMTP) server:' field is empty. Below it, the 'Mail account login details' section contains 'Account Name' and 'Password' fields, both empty, and an 'Authentication Method' dropdown menu set to 'None (anonymous access)'. The 'Sender Information' section has 'Sender Name' set to 'Archive Alert Module' and 'Sender Email' set to 'noreply@archive.com'. A 'Send Test Email' button is located below the sender information. At the bottom of the window are 'Apply', 'OK', and 'Cancel' buttons.

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**

13.10 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 required.

- ❖ 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).

13.11 About On-Screen Messaging


The On-Screen Messaging program can be configured to display via message boxes and system tray notification, as described in [Configuring On-Screen Messaging](#).

It runs on the same computer as the Archive Series software or a connected Windows client. The On-Screen Messaging and Event Monitor are installed automatically on the machine running the Archive Series software at the time of its installation. The On-Screen Messaging program may be installed on a connected Windows client using the Client Utilities installer.



The On-Screen Messaging program connects to the event monitoring service on the computer running the Event Monitor and consequently this must be running. If required, the messaging program may be run simultaneously on multiple clients.

13.12 Configuring On-Screen Messaging

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 computer, 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.

14. Diagnostics & Maintenance

The XenData Archive Series uses the Windows [Event Log](#) to record errors, warnings and informational messages. In addition, it creates [Trace Log](#) messages when an error is encountered.


14.1 Windows Event Log

Whenever the Archive Series 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 and informational messages. An example of an informational message is given below and, in this case, the Archive Series software successfully completed an inventory of the LTO, ODA or object storage at start up.

In general, if the system is not behaving as expected, the Windows Event Log is the first place that you should look.

To Open the Event Log:

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




14.2 System Trace Log

It is sometimes useful to be able to see what is happening internally within the system. The System Trace Log allows you to examine a trace of all actions performed by the system on the LTO, ODA or object storage.

To Open the Trace Log

1. Open the Tiered Storage 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.



Automatic Generation of Trace Files

Whenever the Archive Series 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 at 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.


14.3 Volume Alert State

Certain fault conditions, such as LTO tape errors, ODA disc errors, Object Container or Bucket errors, leave the affected Volume in an "Alert" state. The system does this to protect data by ensuring that it will not attempt to write to Volumes 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 user may decide that it is appropriate to ignore the error and continue to use the affected volume(s). The system provides a mechanism to reset the alert state, as follows:

To Clear the Alert State for a Volume:

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

Note that the Reset Alert option is only available for Volumes that are in the "Alert" state.

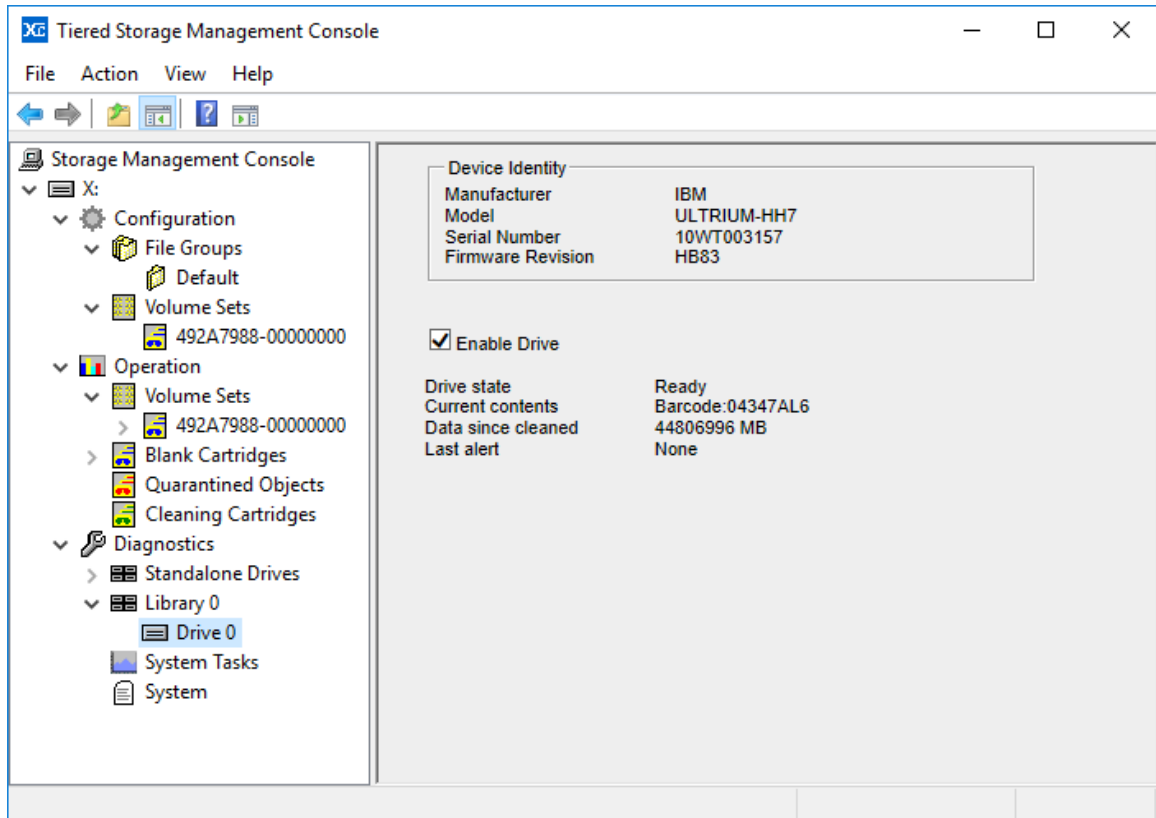


In some cases, it may not be possible to reset the alert state because of the severity of the original condition. In these cases, files written to the Volume prior to the event will be readable. However, additional files cannot be written to the Volume and a new Volume must be allocated to the Volume Set to allow writing of additional files.

14.4 Library and Drive Diagnostic Information

To Open the Diagnostics:

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



14.5 Cleaning LTO Tape Drives

LTO drives, whether internal to a library or stand-alone units, need cleaning from time to time. Cleaning is performed using a cleaning cartridge, which looks much like a data cartridge. The frequency of cleaning varies depending on the cleanliness of the operating environment and of the cartridges. The drives themselves indicate when cleaning is required, via a tape drive alert. For LTO libraries, the Archive Series software will automatically respond to the tape drive alert and will initiate automatic cleaning by moving a cleaning cartridge to the drive. In the case of stand-alone LTO drives, the drive will illuminate an indicator light and the Archive Series software will log a message in the Windows Event Log, send an email alert and provide an on-screen message.


15. System Recovery

If there is a catastrophic failure of the server running the Archive Series software or of the server's disk cache, the system can be [rebuilt from data cartridges or Cloud Containers](#). In a system with one or more LTO or ODA libraries that have a [failed hardware component](#), individual LTO or ODA drives and libraries can be temporarily disabled, allowing the system to continue to function.


15.1 Rebuilding a System from Cartridges or Object Storage

Best practice dictates that the Metadata Backup utility should be used to protect an Archive Series system against catastrophic failure of the server or other disaster. However, Object Storage Containers or Buckets and all the data cartridge formats supported by Archive Series 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 Object Storage Containers, Buckets or the data cartridges using the procedure outlined here.

1. If you have a metadata backup, use the Metadata Restore utility to load the available information into the new system.
 - 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 volumes into the system.
 - Unknown volumes will be recognized by the system and will show in the Tiered Storage Management Console. Volumes written in the LTF5 or ODA formats, Finalized TAR cartridges, Object Storage Containers and Buckets will have Volume Contents Catalogs created on the cache disk.
3. Build the Volume Contents Catalogs for any non-Finalized TAR format tapes, Object Storage Containers or Buckets, using the **Build Missing Catalogs** function shown below. This operation may take several hours because the entire volume 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 volume. 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 volume.



15.2 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. XenData Archive Series software allows the user to selectively [disable hardware](#) while the remainder of the system continues running.

15.2.1 Options in Case of Library or Drive Failure

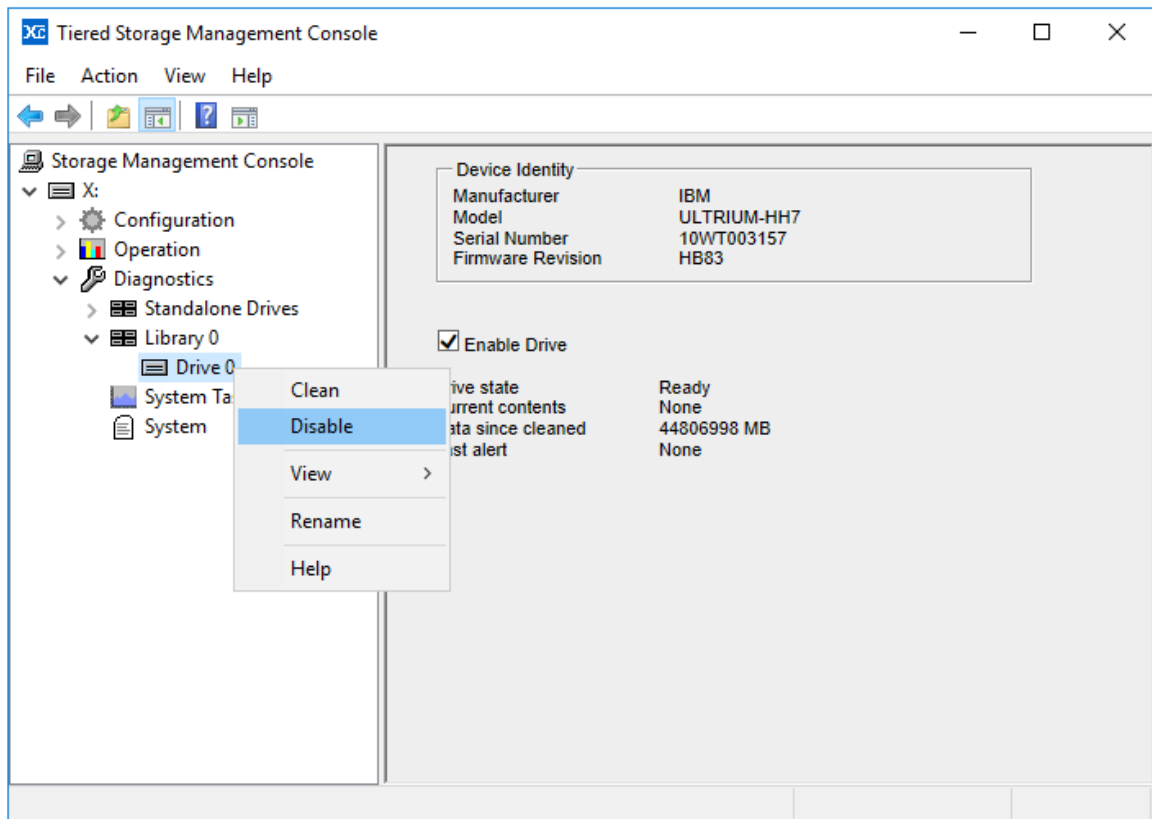
In normal operation, Archive Series software writes files to data cartridge Volumes as they are written to the cache disk. If this is not possible because of a hardware failure, the system will prevent further files from being written. If this behavior is undesirable (perhaps because there is no other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

15.3 Temporarily Disabling LTO or ODA Hardware

To Disable a Drive or Library:

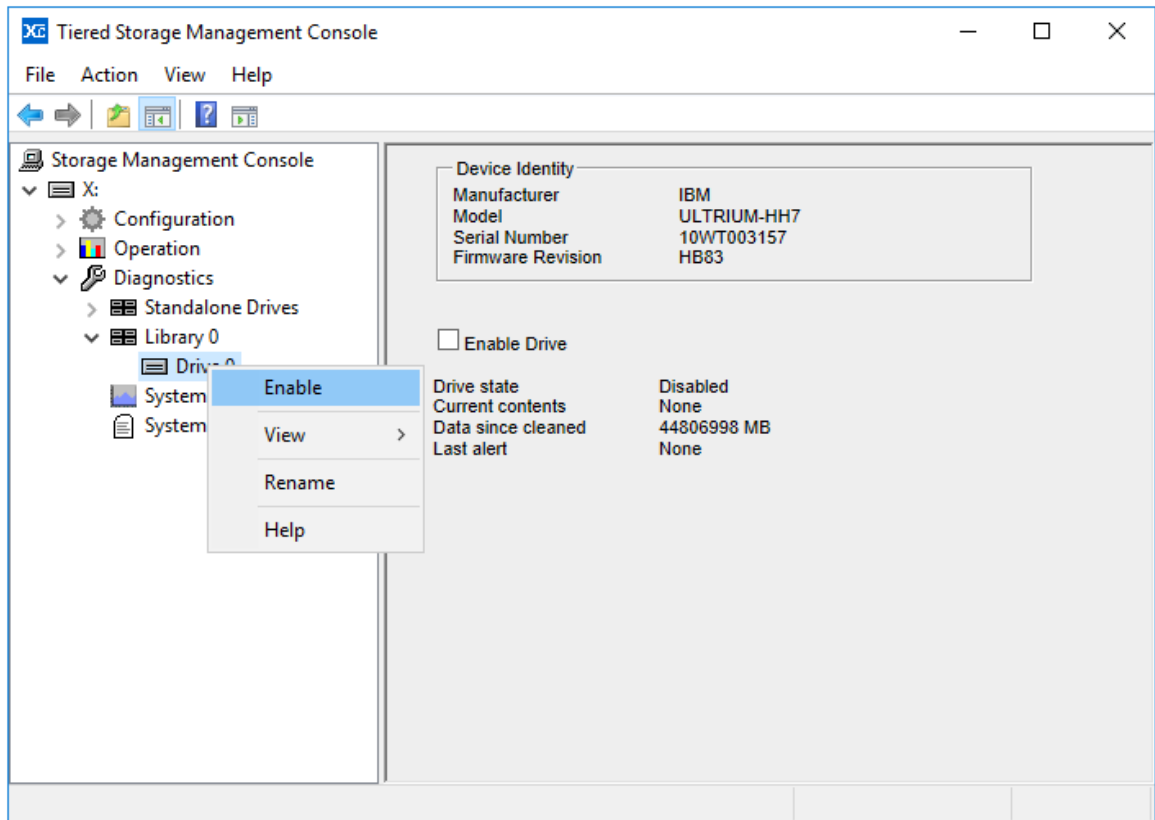
1. Open the Tiered Storage Management Console.

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 Tiered Storage Management Console.
2. Navigate to the **Diagnostics** section.
3. Right-click on the hardware component and select **Enable**.



16. Using Mac Clients

The system running XenData Archive Series software may be used with Mac clients connected via SMB or FTP. However, using Finder to access the XenData file system is not supported.

16.1 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 XenData Archive Series software, 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 Archive Series 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 or by the Report Generator.

16.2 Hidden File Group Policies

In addition to the File Group rules defined using the Tiered Storage 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 Archive Series system will store any file named .DS_Store on the disk cache but will not save the file to LTO, ODA or object storage. This rule overrides all policies defined in the Tiered Storage Management Console.

When a Mac client computer running OS/X writes resource forks and extended attributes to a Windows SMB share (such as a XenData Archive Series 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 tiered 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 Tiered Storage 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 Archive Series system does not flush alternate data streams.

The Object Storage and TAR tape file systems support writing of alternate data streams to Volumes. The other file systems (LTFS and ODA) 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 Volumes, regardless of the format used. The alternate data streams that are treated this way are as follows:

```
:Zone.Identifier  
:AFP_AfpInfo  
:AFP_Afp_Resource  
: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 Tiered Storage Management Console. If the Volume Set specifies writing to a file system that does not support alternate data streams, this will result in an error.

16.3 Disabling Alternate Data Streams

From OS/X version 10.6, Apple uses alternate data streams as the default configuration when writing to a XenData Archive Series SMB share (previously, AppleDouble files were the default). However, the Archive Series 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 XenData Archive Series archive.

17. Client Utilities

The XenData Client Utilities may be installed on a 64 bit Windows 7, Windows 8.1 or Windows 10 client computer connected via a Windows network to the computer running the Archive Series software. There are three utilities that may be installed:

- ❖ [On-Screen Messaging for the Alert Module](#)
- ❖ [File Explorer Extensions](#)
- ❖ [Trace Viewer](#)

17.1 Installing the Client Utilities

1. Download the XenData Client Utilities installer.
2. Run XDClientUtilitiesx64-v.vv.bbbb.xxx.msi (where v.vv is the version number, bbbb is the build number and xxx is a build type).
3. Click 'Next' on the first screen that appears.
4. Click on the 'I accept the terms in the License Agreement' check box, then click 'Next'.
5. For the setup type, click 'Typical' as this is recommended for most users.
6. Click on 'Install'.
7. Once the installation has completed, click on 'Finish'.

17.2 On-Screen Messaging

The On-Screen Messaging program can be configured to display via message boxes and system tray notifications, as described in [Configuring On-Screen Messaging](#).

The On-Screen Messaging program connects to the event monitoring service on the computer running the Event Monitor and consequently this must be running. If required, the messaging program may be run simultaneously on multiple clients.

17.3 File Explorer Extensions

The capabilities of Windows File Explorer on the client computer are extended to provide the following functionality:

- ❖ [Flushing of Files and Folders](#)
- ❖ [Pre-fetching of Files and Folders](#)
- ❖ [Smart Copy and Paste](#)

17.4 Trace File Viewer

Whenever XenData LTO Server Edition 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 at 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 saved in the XenDataLog folder at the root of the system boot drive in a compressed format to make them easier to transmit by email. By installing the Trace Viewer on another Windows computer, you can open and read the contents of a trace file.

17.5 FS Mirror Log Files

When logging is enabled for an FS Mirror or FS Mirror Test task, a log file is created in the XenDataLog folder each time the task runs. These files have the extension .xml and, when viewed on the computer running Archive Series software, a style sheet is launched which displays an easily readable report in Internet Explorer or Microsoft Edge. By installing the Trace Viewer on another Windows computer, you extend the ability to read an FS Mirror or FS Mirror Test log file to that computer. Note that the log file must be transferred to that computer as IE and Edge security prevent operation over a network.

18. Glossary

Activation Code An Activation Code is required to run the Archive Series software and enables the chosen configuration. Separate Activation Codes are required to enable Multi-Site Sync, FS Mirror and Alert Module functionality. The Archive Series License Administration utility is used to apply activation codes to a system.

Alert Module It provides email and on-screen alerts that are tailored to the needs of systems administrators and 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'.

Amazon Web Services is Amazon's public cloud computing platform. It provides a comprehensive range of services, including computing, analytics and data storage.

API is an acronym for 'Application Program Interface'. XenData APIs are available to software developers to tightly integrate their applications with the Archive Series 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 '._'.

Azure is Microsoft's public cloud computing platform. It provides a comprehensive range of services, including computing, analytics and data storage.

Blank Cartridge Set is applicable to XenData software that manages LTO or ODA cartridges. It is the set of data cartridges shown in the Management Console which consist of new (unused) cartridges or rewritable cartridges that have been reformatted.

Blob Storage is Microsoft's name for Azure Object Storage. The name is derived from 'Binary Large Object'. A Blob is a stored Object and all Blobs are grouped in Containers.

Buckets are object repositories, used by the Amazon and Wasabi implementations of S3 to hold and organize individual objects.

Cache Disk is the magnetic or solid state disk volume under control of the Archive Series software. It is also termed 'managed disk'.

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.

Tiered Storage Management Console Used to configure all File Group, Volume Set settings and to view diagnostic information about the system

Container An Azure Container represents a grouping of Blobs.

Contents Catalog The Archive Series software creates a Contents Catalog for each Volume that it creates. This is stored on the disk cache as a hidden file.

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

Event Log See Windows Event Log.

File-Folder Interface This is a term used in this User Manual to refer to the file system contained in the logical drive letter that is managed by the Archive Series software.

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 storage of very large files. This option is not available for the default settings of the Cloud File Gateway.

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 and have the same disk retention policy). Files are assigned to a File Group on the basis of their names.

Finalization Process that writes a contents catalog for a Volume to the LTO, ODA or object storage. After the Volume has been Finalized, no additional files may be written to that Volume.

Flushing Files are flushed when they are removed to free space on a storage device. The Archive Series software can be configured to automatically flush files from the disk cache once they are securely stored on LTO, ODA or object storage. After flushing, the file remains visible at the same location in the file-folder interface, however is displayed as 'offline'. When the offline file is read, it is restored automatically from the LTO, ODA or object storage.

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.

Generic S3 is any S3 implementation that does not have its own implementation officially named and supported within XenData.

History Explorer within Windows File Explorer is used to obtain the version history and status of any file, including deleted and renamed files.

HTTP, or HyperText Transfer Protocol, is a transport protocol which is the foundation of almost all data exchange on the Web.

HTTPS, or HyperText Transfer Protocol Secure, is an extension to HTTP which encrypts data while in transit. Furthermore, it uses a certificate issued by a certification authority to verify that the connection is legitimate. The certificate is known as an SSL Certificate.

LTFS An acronym for 'Linear Tape File System'. It is a tape cartridge format supported by the LTO Edition of XenData Archive Series software. It is the most popular format for archival applications and defines how file data and file system metadata are written to tape cartridges. It allows cartridge interchange between LTO systems from different manufacturers that support LTFS. It is applicable to rewritable LTO cartridges but cannot be used with WORM cartridges.

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

Managed Disk is the magnetic or solid state disk volume under control of the XenData Archive Series software. It is also termed 'cache disk'.

MMC An acronym for 'Microsoft Management Console'. It can be used to create, save, and open administrative tools that manage the hardware, software and network components of a Windows system. The Tiered Storage Management Console is an example of such a tool.

Multi-Site Sync is a XenData service used with two or more Cloud File Gateway instances that share access to cloud object storage. The service shares the file-folder structure created by each gateway with all gateways.

Named Streams See Alternate Data Streams.

NFS An acronym for 'Network File System'. It is the standard protocol used by Unix and Linux computers to communicate over a network.

NTFS Microsoft's 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 XenData Archive Series managed disk must be formatted with NTFS.

Object Store is a generic term covering object-based storage mediums.

Object-Based Storage is a computer data storage architecture that manages data as objects, as opposed to file system architectures which managed data as a file hierarchy, and block storage which manages data as blocks within sectors and tracks.

Offline File Attribute A file attribute bit defined by Microsoft. XenData Archive Series software sets the offline file attribute bit to identify files that have been flushed from the managed disk.

ODA See Optical Disc Archive.

Optical Disc Archive is a storage technology that was introduced by Sony. In this documentation it is also termed 'ODA'. It uses removable cartridges, where each ODA cartridge holds 11 or 12 optical discs. Each of the internal optical discs is similar to a Blu-ray disc.

Petabyte 1024 terabytes. It is abbreviated to PB.

Quarantined Volume In the case of the Cloud File Gateway, it is an Object Storage Container or Bucket having Objects that are not available to the XenData software. In the case of LTO or ODA, it is a location in the Management Console 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.

S3, or Simple Storage Service, is a service developed by Amazon Web Services (AWS) that provides object storage through a web services interface. The S3 interface has been adopted by many other companies, including XenData, as an interface for both private and public cloud object storage.

SMB See CIFS.

SSL is a security technology for establishing an encrypted link between a server and a client, such as a website and a browser. An SSL certificate is a digital certificate that authenticates the server's identity and enables an encrypted connection.

State File An XML file that contains configuration settings for the Tiered Storage Management Console including File Group and Volume Set configuration settings.

TAR is a term derived from 'Tape ARchive' and is a tape cartridge format supported by the LTO Edition of Archive Series software for both rewritable and WORM cartridges.

Terabyte 1024 gigabytes. It is abbreviated to TB.

Ultrium See LTO.

Volume For the Cloud File Gateway, it is an Object Storage Container or Bucket. For ODA, it is an ODA cartridge. For LTO, it is set of replicated tape cartridges.

Volume Set A set of one or more Volumes which store files from designated File Groups.

Wasabi is a public cloud storage platform, which uses a modified implementation of Amazon's S3.

Windows Event Log The XenData Archive Series software provides a comprehensive array of warnings and informational messages which are logged in the Windows Event Log. In general, if the system is not behaving as expected, the Windows Event Log is the first place that you should look.

WORM is an acronym for 'Write Once Read Many'. WORM tape and optical cartridges cannot be reformatted and after data is written to a WORM cartridge, it cannot be changed.