

XenData Cloud File Gateway User Manual

®

Including Multi-Sync and S3 Server Interface Options

Version 7.25.4270.100

1	About The Software	6
2	Introduction	8
2.1	Overview	9
2.2	Writing Files to the System	11
2.3	Reading Files from the System	13
2.4	Antivirus Software Compatibility	13
2.5	Synchronizing and Importing Files Written to Object Storage by Another System	14
2.6	Overview of the S3 Server Interface	14
2.7	Multi-Site Sync Service	16
2.7.1	Handling of File Versions, Deletions, Overwriting and Renaming	17
2.7.2	How Files are Stored and Synchronized	18
2.7.3	About Application Specific Optimizations	18
2.7.4	Microsoft Office Optimizations	19
2.7.5	Adobe Premiere Pro Optimization	20
2.8	CFG Functionality - In Summary	21
3	File Operations, Security and Connectivity	24
3.1	Supported File and Folder Operations	25
3.2	Unsupported Rename Operations	25
3.3	About Stub Files and the Offline Attribute	25
3.4	Handling of Alternate Data Streams	26
3.5	Supported Network Protocols	27
3.6	Cloud Gateway Free Space Reporting	27
3.7	File Security	27
4	Concepts	28
4.1	About File Groups	29
4.2	About Objects, Volumes and Volume Sets	29
4.3	About Volume Catalogs	29
4.4	About Volume Finalization	30
4.5	About Pending Write Mode	30
5	Administering the System	31
5.1	Cloud Gateway Management Console	32
5.2	Azure Storage Accounts	33

5.2.1	Adding Azure Storage Account Access	33
5.2.2	Adding Azure Key Vault Access	35
5.2.3	Configuring a Storage Account	36
5.3	Amazon S3 Endpoints	37
5.3.1	Adding Amazon S3 Account Access	37
5.3.2	Configuring an Amazon S3 Account	38
5.4	Wasabi S3 Endpoints	39
5.4.1	Adding Wasabi S3 Account Access	39
5.4.2	Configuring a Wasabi S3 Account	40
5.5	Multi-Site Sync Service	41
5.5.1	Adding the Multi-Site Sync Service	41
5.5.2	Adding Cosmos DB Account Access	41
5.6	Configuring Generic S3 Endpoints	42
5.6.1	Adding Generic S3 Account Access	42
5.6.2	Configuring a Generic S3 Account	44
5.7	Volume Sets	44
5.7.1	Creating a New Volume Set	45
5.7.2	Renaming a Volume Set	45
5.7.3	Configuring a Volume Set for Cloud Storage	45
5.7.4	Scanning for Object Storage Containers Created by Other Systems	47
5.7.5	Adding a Volume	48
5.7.6	Deleting a Volume Set	48
5.7.7	Deleting a Volume	49
5.7.8	Rebuilding Volume Contents Catalogs	49
5.7.9	Import Folder Structure	49
5.7.10	Import Data	50
5.7.11	Obtaining Volume Statistics	50
5.7.12	Write Protecting a Volume	51
5.7.13	Finalizing Volumes	52
5.8	File Groups	52
5.8.1	Creating a New File Group	52
5.8.2	Renaming a File Group	53
5.8.3	Changing the Order of File Groups	53
5.8.4	Allocating Files to a File Group	53
5.8.5	Selecting a Volume Set for a File Group	55
5.8.6	Selecting Disk Retention Rules	55
5.8.7	Changing Disk Retention Rules	57
5.8.8	File Group Advanced Options	57
5.9	S3 Server Interface	59
5.9.1	Adding the S3 Server Interface	59

5.9.2	Configuring the S3 Server Interface	60
5.9.2.1	Host	60
5.9.2.2	Access Keys	60
5.9.2.3	Buckets	62
6	File Explorer Plug-In	63
6.1	Flushing of Files and Folders	64
6.2	Pre-fetching of Files and Folders	64
6.3	Smart Copy and Paste	65
6.4	Enhanced Properties	66
6.4.1	Obtaining Enhanced Properties	66
6.4.2	Changing the Object Storage Tier for a File	66
6.5	Volume View	67
6.6	History Explorer for Cloud File Gateway	68
7	Metadata Backup	69
7.1	Starting Metadata Backup	70
7.2	Selecting Backup or Restore	70
7.3	Making a Predefined Backup	71
7.4	Making a Custom Backup	73
7.5	Restoring a Backup	77
8	Scheduler	81
8.1	Starting the Scheduler	82
8.2	Adding a Task	82
8.3	The Scheduler Status Display	83
8.4	Editing and Deleting Tasks	84
8.5	Starting and Stopping Tasks	85
8.6	Scheduling Metadata Backup	86
8.7	Scheduling Deferred Write	87
8.8	Scheduling File System Mirror	88
8.9	Scheduling File System Mirror Reporting Run	92
9	Reports	95
9.1	Starting the Report Generator	96
9.2	Creating, Saving and Restoring Reports	96
9.3	File Search Report	97
9.3.1	Interpreting a File Search Report	99

9.4	UnArchived Files Report	100
9.4.1	Interpreting an UnArchived Files Report	101
9.5	Volume Contents Report	102
9.5.1	Interpreting a Volume Contents Report	103
10	Alert Module	105
10.1	About the Event Monitor	106
10.2	Configuring the Event Monitor	106
10.3	About Event Categories	107
10.4	Configuring Event Categories	108
10.5	About Recipient Groups	109
10.6	Configuring Recipient Groups	109
10.7	About the Email Server	110
10.8	Configuring the Email Server	110
10.9	Error Reporting	112
10.10	About On-Screen Messaging	113
10.11	Configuring On-Screen Messaging	113
11	Diagnostics & Maintenance	116
11.1	Windows Event Log	117
11.2	System Trace Log	117
11.3	Active Files Display	118
11.4	System Tasks	121
12	Client Utilities	123
12.1	Installation Prerequisites	124
12.2	Installing the Client Utilities	124
12.3	On-Screen Messaging	124
12.4	File Explorer Plug-In	125
12.5	Trace File Viewer	125
12.6	FS Mirror Log Files	125
13	Glossary	126

1. About The Software

Cloud File Gateway Software

Including Multi-Sync and S3 Server Interface Options

Version 7.25.4270.100

Copyright © 2001-2023 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, the Cloud File Gateway Extension which adds the ability to archive to object storage and the Multi-Site Sync Extension that creates a synchronized file system across multiple gateway instances.

This user documentation describes the Cloud File Gateway Edition and its optional extensions. It does not support LTO data tape or Optical Disc Archive cartridges.

2.1 Overview

XenData Cloud File Gateway software allows file-based applications to use Object Storage. It runs on a Windows platform which can be a physical on-premise computer, an on-premise virtual machine (VM) or a VM running in the cloud. Supported operating systems include Windows 10 Pro, Windows Server 2019 and Windows Server 2016. For an up-to-date list, please refer to the support section of the XenData web site at www.xendata.com/support. Each gateway manages a local disk volume, often termed the cache disk, to provide enhanced performance and provides highly granular tiering policies for optimizing file retention on disk, as described later in this section.

A single computer or VM running the Cloud File Gateway is illustrated below.



Multiple Cloud File Gateway instances may be synchronized using XenData's Multi-Site Sync service. This provides a global file system accessible from each gateway and sharable to locally attached networks. 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 and is described further in [this section](#).

The Cloud File Gateway is optimized for large unstructured data files including those containing video. Restoring partial files via the gateway is supported which is particularly useful when only a portion of a large video file is required to create a video clip. The gateway supports streaming of video files and allows scrubbing along the timeline of a video player.

Each Cloud File Gateway allows multiple file-based applications to simultaneously use object storage which is accessed locally or via one or more network shares using SMB, NFS or FTP network protocols. It integrates fully with the Microsoft security model based on Active Directory. Consequently, it can be easily added to an existing Windows domain or workgroup.

The gateway supports a range of Object Storage types including Amazon Web Services S3, Azure Blob Storage and Wasabi S3. For an up-to-date list, please refer to the XenData web site at www.xendata.com. The connection to Object Storage uses HTTPS. A single Cloud File Gateway instance supports simultaneous connection to multiple storage accounts including selected endpoint locations from a single Object Storage provider as well as storage accounts and endpoints from different providers. The gateway may be configured to replicate files across different endpoints.

The Cloud File Gateway writes to Object Storage in a standard way and, in its default configuration, each file is written to one object. The gateway is compatible with many other applications. For example, many utilities provided by public cloud providers can restore files written with the gateway. Also, the gateway can index and manage objects written by other standard applications.

A single instance of the Cloud File Gateway is highly scalable, supporting object storage of unlimited capacity with up to 2 billion managed objects. The gateway will manage a local disk volume up to 256 TB. It takes advantage of high bandwidth connections between the gateway and the object storage, employing multi-part upload and download that is optimized for large video files and other unstructured data files.

Files are written and read via the gateway just as you would to a Windows share on a file server or to a local logical drive on the computer or VM that hosts the XenData gateway software. XenData Cloud File Gateway tiering policies keep frequently accessed files on the managed disk volume. When the gateway software is installed on an on-premise computer, this minimizes cloud access costs and Internet bandwidth usage. Different tiering policies may be applied to specified folders and files, allowing a great deal of optimization. Options include storing file instances on disk only, Object Storage only and both disk and Object Storage. A further option is to initially store files on both disk and Object Storage and configure a disk retention period after which a file is available from Object Storage only. When files are stored only on Object Storage, they continue to appear in the file-folder interface providing seamless access from file-based applications.

Object Storage accounts store objects in containers or buckets. These are also termed Volumes in this guide. The Cloud File Gateway automatically creates Volumes and writes objects to them.


When a Volume has 1 million files, the gateway identifies it as full and then jumps to a new Volume to write additional objects. When using the Cloud File Gateway to write and read files to Object Storage, the management of Volumes is transparent to the user.

The Cloud File Gateway may be upgraded to include an S3 Server interface which provides access to the gateway including the managed disk volume. The S3 Server interface allows files to be written to and read from the Cloud File Gateway using HTTP or HTTPS. It does not affect access to the gateway as a local logical drive letter or over the local network using SMB, FTP or NFS. The S3 Server interface is described further in [Overview of S3 Server Interface](#).


The Cloud File Gateway may be upgraded to include a file system synchronization utility called FS Mirror. This allows mirroring between any file-folder structure accessible to the Windows machine running the gateway. It uses end-to-end checksum protection to verify file integrity. A common use is mirroring selected folders on a shared disk volume on the network connected to the gateway to object storage and vice-versa. Another use-case is to map one folder in the gateway to one endpoint and another folder to a different endpoint; then use FS Mirror to replicate the folders across the endpoints. FS Mirror runs as a scheduled task and is described further in [Scheduling File System Mirror](#).

2.2 Writing Files to the System


For each computer or VM that is running the Cloud File Gateway, the managed file system 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 gateway. When files are written to the system, they are always first written to the cache disk as illustrated here.



After a file is written to disk, it will be copied to the Object Storage – that is if the policy is configured to do this. Immediately after the Object Storage copy is created, there are two instances of the file, as illustrated below.



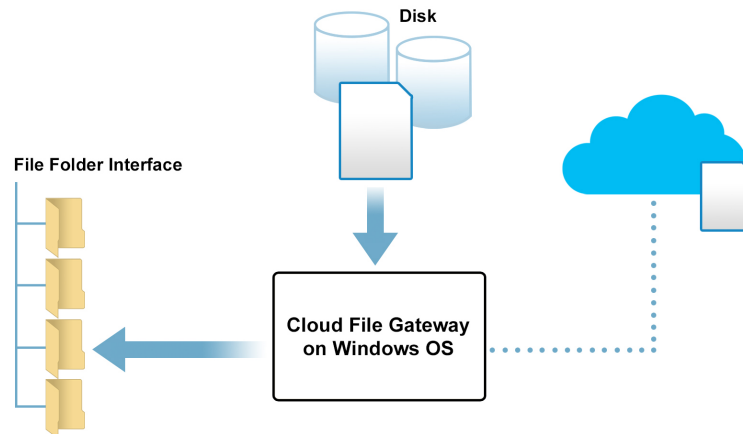
If the storage policy is configured to retain only the Object Storage copy, then the instance of the file on the cache disk will be converted to a stub file. This process of conversion from a full file to a stub file is called flushing. The sparse file has all the characteristics of the original file, such as size and modification date, except there is one file attribute that changes. This is the offline attribute which indicates that the file is no longer available from disk. The flushing operation frees up space on the managed disk because the stub file takes only a few KBytes.




Conversion to a stub file may be scheduled to occur immediately after the file has been successfully written to Object Storage. Alternatively it may be scheduled to occur a defined time after the file was written or last read. Or another option is that some files may be retained on disk cache permanently, as well as being written to Object Storage. The rules that determine how long the full file will be retained on disk are defined by using the Cloud Gateway Management Console. The retention rules may be different for different file types or for different folders.

2.3 Reading Files from the System

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



When the file is a stub file representation on the managed cache disk and the full file is on Object Storage, it is restored directly and automatically from the Object Storage.



2.4 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.

2.5 Synchronizing and 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](#)

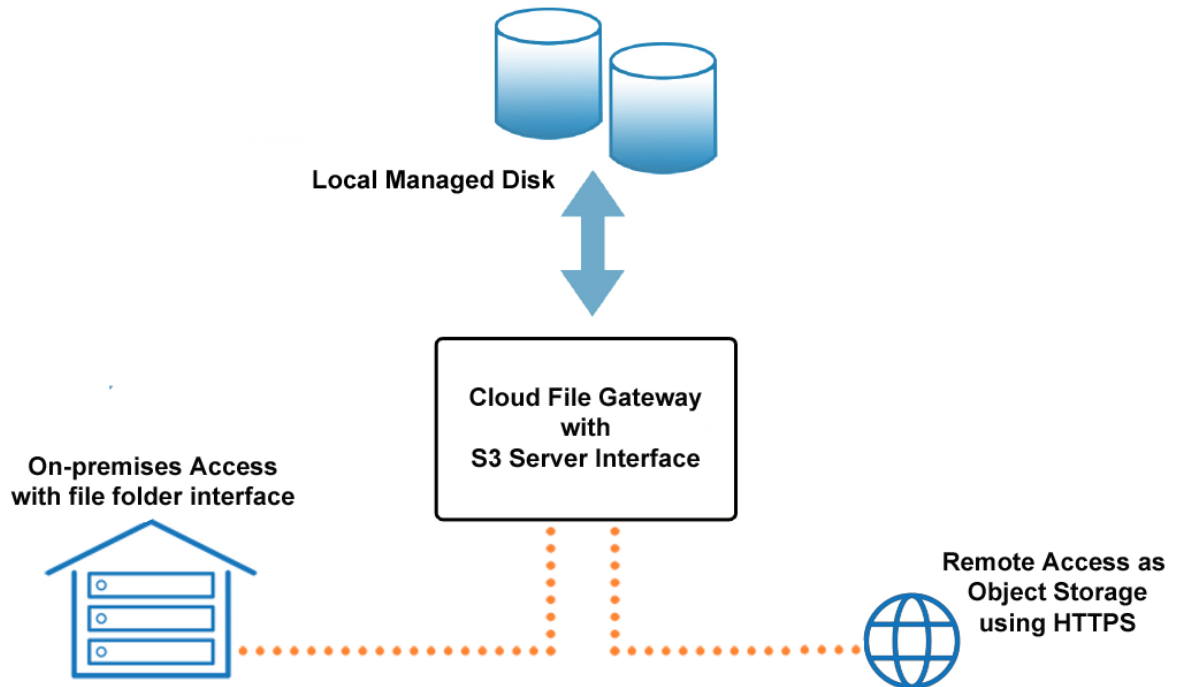
A scheduled PowerShell script may be used to perform these steps automatically.

If the Object Storage is Azure or AWS, you can forgo the above steps thanks to the Change Feed Extension. With the Change Feed Extension, file creations, updates and deletes on a connected external container are quickly replicated on your XenData system, without the need for manual intervention. The Change Feed Extension requires a separate XenData service, which can be found in the Cloud File Gateway, or Cloud File Extension installers. The Change feed works on externally created containers or buckets, and has to be enabled on the Storage Account; there are different procedures for configuring this depending the Object Storage provider. More details can be found in section 2.2 Installing Multi-Site-Sync or Change Feed Extension in the Cloud File Gateway Installation manual.

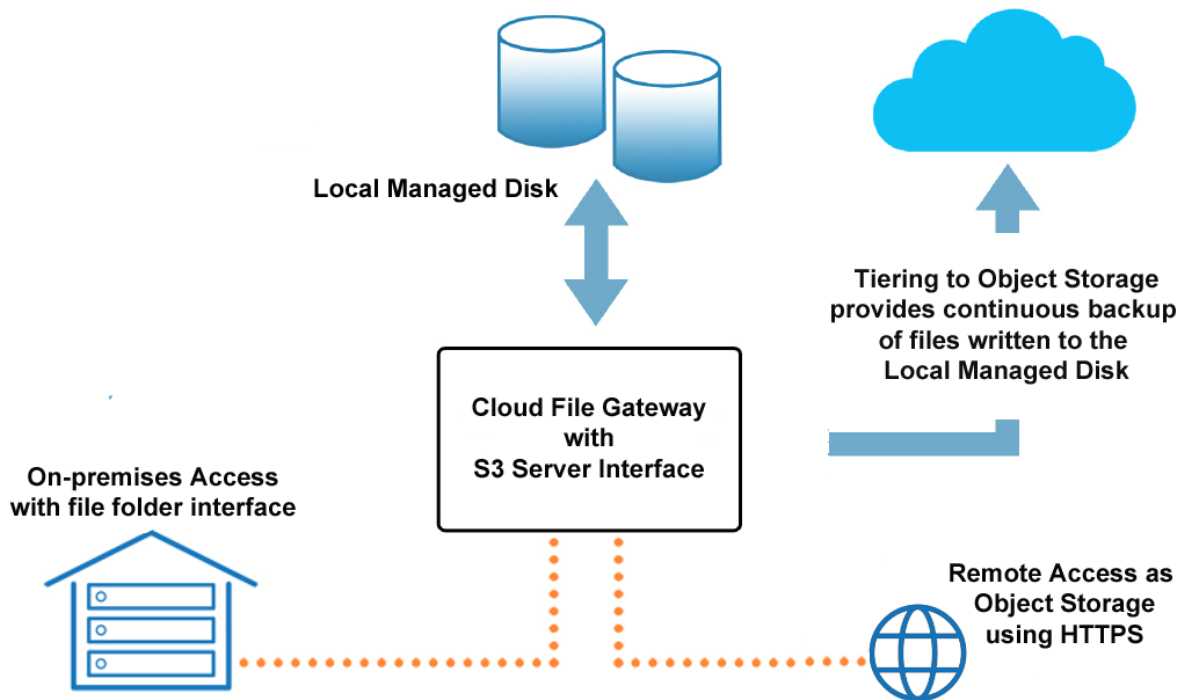
2.6 Overview of the S3 Server Interface

The optional S3 Server Interface provides access to the Cloud File Gateway using HTTP or HTTPS.

The S3 Server Interface is often used to simply write to and read from the managed disk volume as object storage. This may be accessed remotely using secure HTTPS. The managed disk volume continues to be accessible locally as a logical drive letter and over the local network using SMB, FTP or NFS. In this case, only the S3 Server Interface is licensed and the system is used to provide disk-based object storage, as shown below.




The functionality of the system may be further extended by licensing the gateway to tier files written to the local disk volume to object storage. This tiered object storage may be used to backup the content stored on the local managed disk, as illustrated below.



2.7 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.

2.7.1 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.

2.7.2 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.

2.7.3 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.

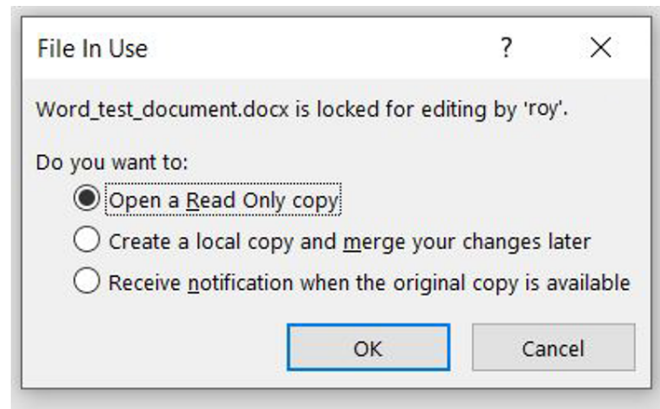
2.7.4 **Microsoft Office Optimizations**

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.




2.7.5 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.

2.8 CFG Functionality - In Summary

The standard Cloud File Gateway license provides the following functionality.

- ❖ File Access to Object Storage - Allows existing file-based applications to use massively scalable, shared cloud Object Storage without modification to support Object Storage APIs.
- ❖ The gateway runs on a Windows VM in the cloud, on a Windows VM on-premises or on a physical Windows server or Windows 10 computer.
- ❖ Multi-Cloud - The gateway simultaneously writes to and reads from a range of Object Storage types including Amazon Web Services S3, Azure Blob Storage and Wasabi S3.
- ❖ Supports Import of Objects from 3rd Party Applications - containers or buckets created by 3rd party applications such as Azure Storage Explorer, AzCopy, AWS CLI or Wasabi Explorer may be imported into the gateway file system.

- ❖ Disk Cache - Each computer or VM running the gateway software has a managed disk which may be configured to cache frequently accessed files.
- ❖ Tailored Disk and Object Storage Policies - Tiering policies determine which files are written to the Object Storage, disk cache or both. The user can set policy rules for various file types and folders.
- ❖ Timed disk retention rules may also be applied to files written to the Object Storage, determining how long files are cached on disk after written or last read.
- ❖ Supports Pre-Fetch and Flush of Files - The pre-fetch operation creates an instance of a file on the disk cache copying it from Object Storage. The flush operation removes a cached file from the disk, replacing it with a stub file to free space on the disk.
- ❖ Check-sum Verification – Automated end-to-end check-sums are used to verify that files are correctly written to Object Storage.
- ❖ Scheduler Optimizes Internet bandwidth - Time windows may be scheduled so applications can write to the disk cache while postponing a copy being made to Object Storage. When the gateway is installed on an on-premise computer, it allows Internet bandwidth to be optimized when in high demand.
- ❖ Partial File Restore (PFR) – Allows restore of a portion of a file, avoiding the need to restore the entire file.
- ❖ Metadata Backup and Restore – A file system metadata backup and restore utility is provided. Benefit: provides rapid system restore, in case of rebuild after failure of the disk cache volume.
- ❖ Alert Module – A software module is included which provides e-mail and on-screen alerts.
- ❖ 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.
- ❖ Industry Standard File Security - The Cloud File Gateway runs on a Windows operating system and integrates fully with the Microsoft Windows security model based on Active Directory.

The standard Cloud File Gateway license may be upgraded to provide the following additional functionality.

- ❖ Multi-Site Sync - Multiple Cloud File Gateway instances may be synchronized using XenData's Multi-Site Sync service. This provides a global file system accessible from each gateway and sharable to locally attached networks. 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.

- ❖ Local Sync - The Cloud File Gateway may be upgraded to include a file system synchronization utility called FS Mirror which allows mirroring between any file-folder structure accessible to the Windows machine running the gateway. A common use is mirroring selected folders on a shared disk volume on the network connected to the gateway to object storage and vice-versa. Another use-case is to map one folder in the gateway to one endpoint and another folder to a different endpoint; then use FS Mirror to replicate the folders across the endpoints. FS Mirror runs as a scheduled task and is described further in [Scheduling FS Mirror](#).
- ❖ S3 Server Interface – This upgrade allows files to be written to and read from the Cloud File Gateway using HTTP or HTTPS.

3. File Operations, Security and Connectivity

The XenData Cloud File Gateway 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.

3.1 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.

3.2 Unsupported Rename Operations

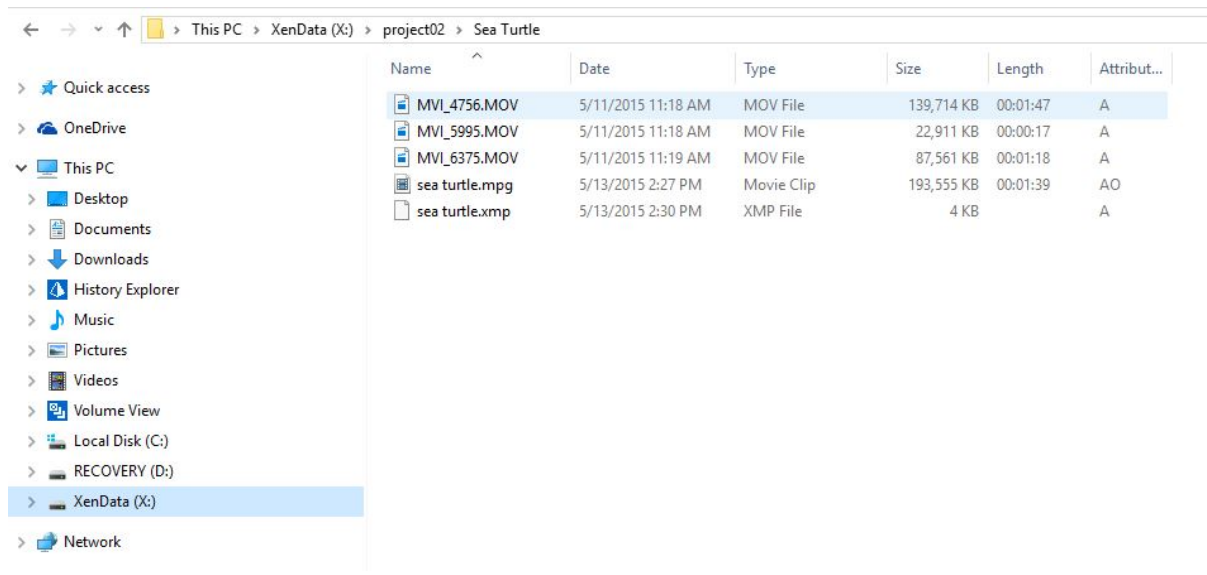
The Cloud File Gateway software does not support renaming folders after a file has been added to that folder.

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

3.3 About Stub Files and the Offline Attribute

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.

3.4 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 Cloud File Gateway software. These alternate data streams contain application-specific file metadata and/or Finder display layout information. The Cloud File Gateway software preserves Finder display information on the cache disk but does not write it to the 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 Cloud File Gateway software preserves this information on the disk cache but does not attempt to write it to Object Storage.
- ❖ Other types of application-specific alternate data streams will be written to Object Storage in addition to the disk.

3.5 Supported Network Protocols

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

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

3.6 Cloud Gateway Free Space Reporting

The total space reported for the logical disk volume managed by the Cloud File Gateway is the licensed capacity of the Object Storage; the free space is reported as the difference between the total space and the used space.

3.7 File Security

The Cloud File Gateway 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.

4. Concepts

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

4.1 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 Cloud File Gateway 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.

4.2 About Objects, Volumes and Volume Sets

Microsoft Azure and S3 object storage use different terminology for objects and how they are grouped.

An Azure object is termed a Blob, whereas it is simply termed an Object by S3 storage providers. In the case of Azure, the XenData Cloud File Gateway in its default configuration writes each file to one Blob; and in the case of S3, each file is written to one Object.

Within their storage accounts, Azure Blobs and S3 Objects are grouped in Containers and Buckets, respectively. An Azure Container and an S3 Bucket is termed a Volume by XenData. 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, 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.

If an Azure Container or S3 Bucket is not usable by the system, it will be identified as a Quarantined Volume in the Cloud Gateway Management Console.

4.3 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.

Volumes may be [Finalized](#) which prevents additional files being written to that Volume. The Finalization process writes an instance of the Volume Catalog to an Azure Container or S3 Bucket dedicated to storage of the Volume Catalogs.

4.4 About Volume Finalization

Volume finalization prevents additional files being written to that Volume. Finalization is performed automatically when a Volume becomes full and may be initiated manually from the Cloud Gateway Management Console. The Finalization process writes the [Volume Catalog](#) to a separate Azure Blob Storage Container or S3 Bucket.

4.5 About Pending Write Mode

In normal operation, the Cloud File Gateway software writes files to the designated 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 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.

5. Administering the System

The main interface for managing the system is the Cloud Gateway 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.

5.1 Cloud Gateway Management Console

The Cloud Gateway 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 Cloud Gateway Management Console


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

To Navigate the Cloud Gateway 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.

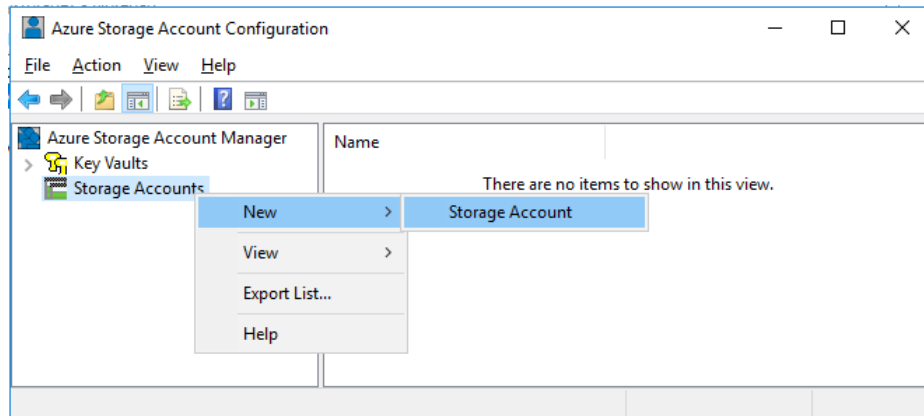


5.2 Azure Storage Accounts


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

5.2.1 Adding Azure Storage Account Access

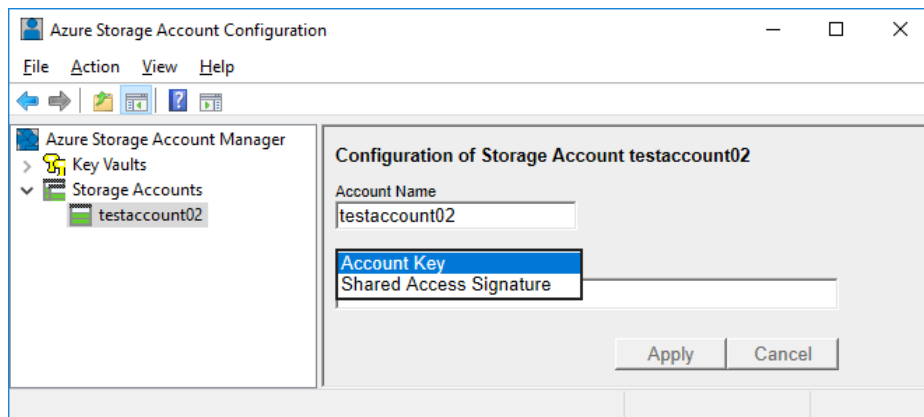
1. Launch the Azure Storage Account Configuration utility as follows:
 1. Click the Windows Start icon.
 2. Open the XenData program group
 3. Click the **Azure Storage Account Configuration** entry in the list.
2. Right-click on 'Storage Accounts'; select 'New' and then 'Storage Account'.



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'.

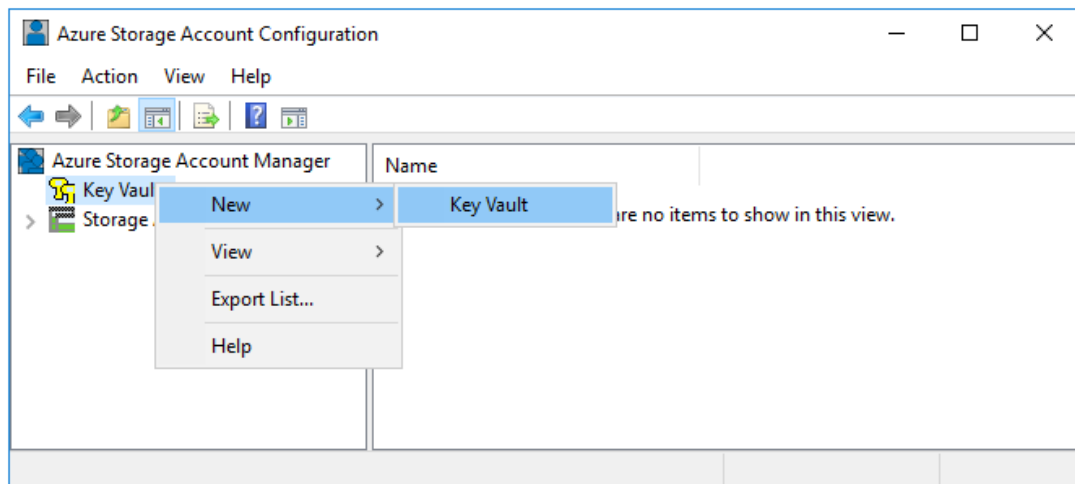


5. Reboot the computer.


5.2.2 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 Cloud File Gateway access to the accounts controlled by a key vault by giving it credentials for the key vault.

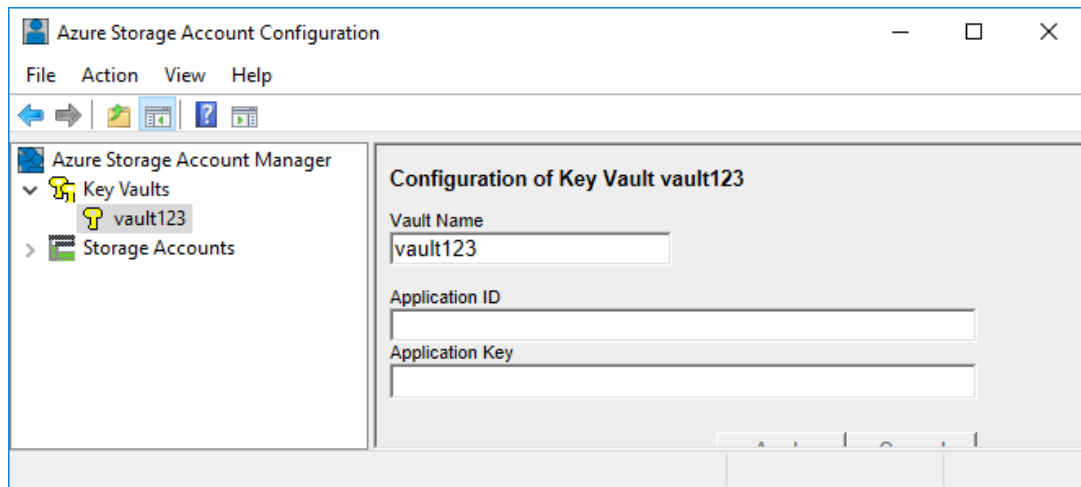
1. Launch the Azure Storage Account Configuration utility as follows:
 1. Click the Windows Start icon.
 2. Open the XenData program group.
 3. Click **Azure Storage Account Configuration** in the list.
2. Right-click on 'Key Vaults'; select 'New' and then 'Key Vault'.



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



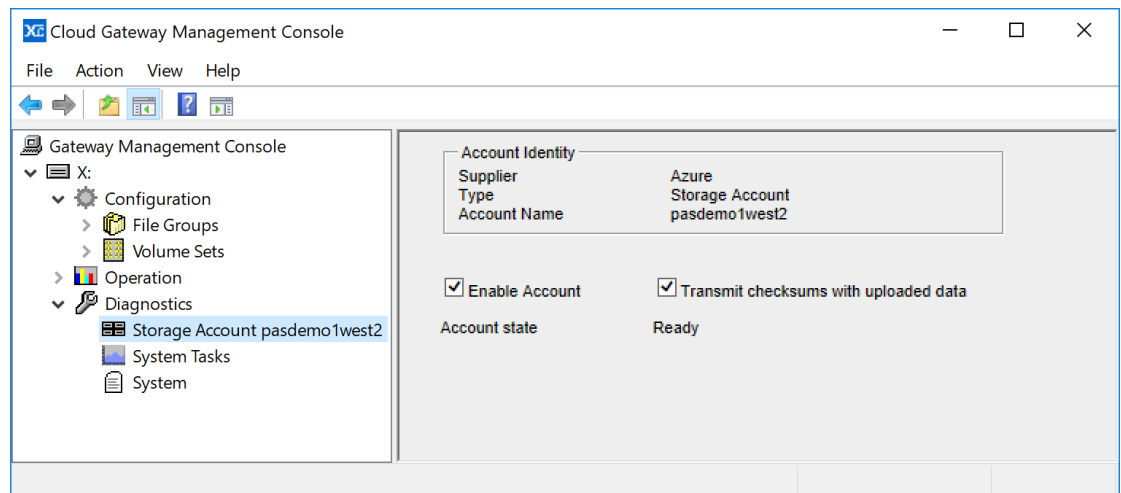
4. Left-click on the key vault name shown under 'Key Vaults', enter the settings for the key vault, then click 'Apply'.



5. Reboot the computer.

5.2.3 Configuring a Storage Account

1. Expand the **Diagnostics** section in the left pane of the Cloud Gateway 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.

5.3 Amazon S3 Endpoints


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

5.3.1 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'



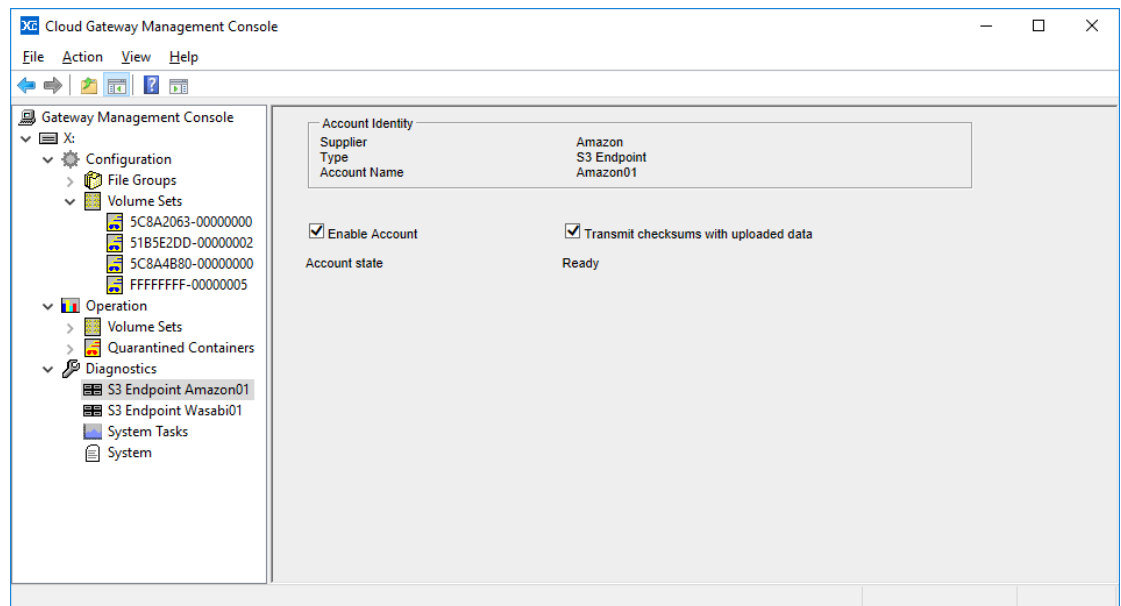
4. Left-click on the endpoint name shown under 'S3 Endpoints', then enter the 'Access Key ID' and 'Secret Key' from your S3 account. Once you've added the keys, select the region you wish your Buckets to be created in from the drop down, and click 'Apply'.



5. Reboot the computer.

5.3.2 Configuring an Amazon S3 Account

1. Expand the **Diagnostics** section in the left pane of the Cloud Gateway 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.

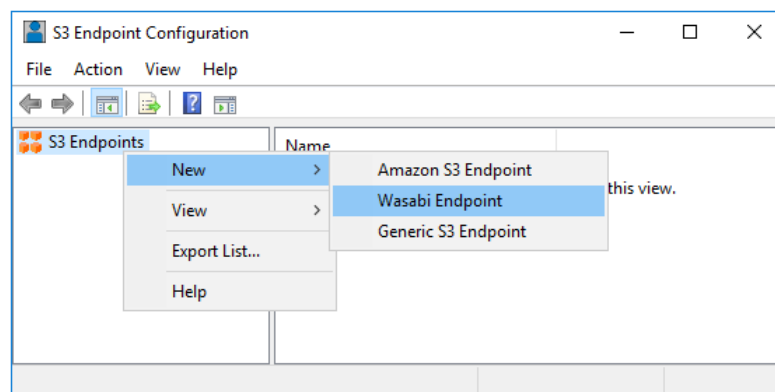
- ❖ **Transmit checksums with uploaded data.** By enabling this option, checksums are transmitted when data is uploaded to the endpoint and are used for data verification purposes.

5.4 Wasabi S3 Endpoints


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

5.4.1 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

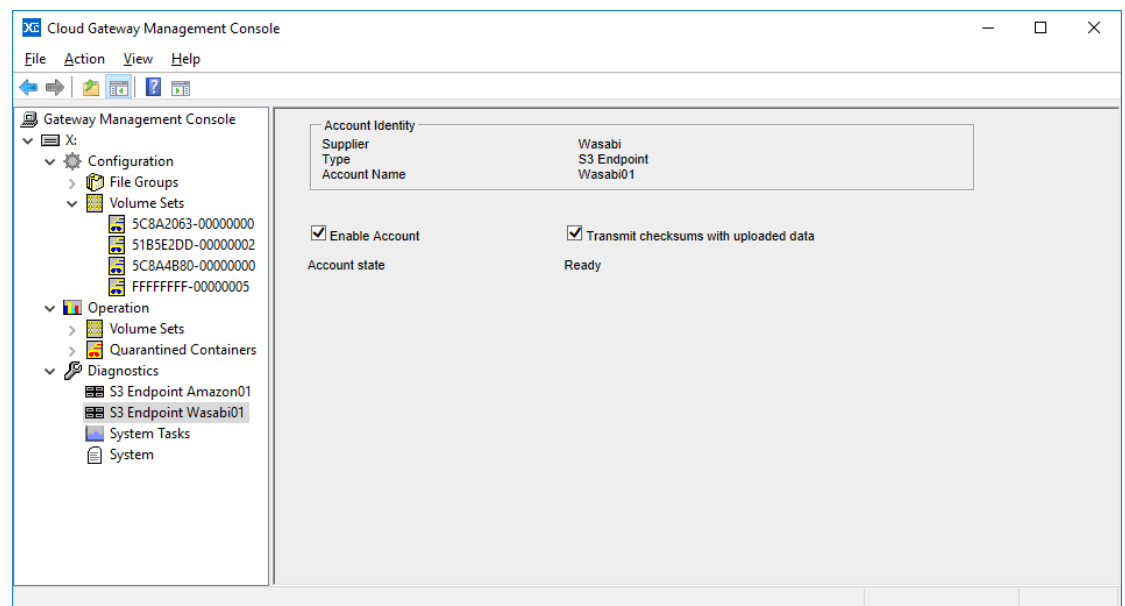
streams of data to the S3 bucket for increased performance. This option has a higher risk of write failures on some internet connections, so should be tested on your connection before production use.



5. Reboot the computer.

5.4.2 Configuring a Wasabi S3 Account

1. Expand the **Diagnostics** section in the left pane of the Cloud Gateway 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.

5.5 Multi-Site Sync Service

5.5.1 Adding the Multi-Site Sync Service

The [Multi-Site Sync](#) functionality for Cloud Object Storage 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:

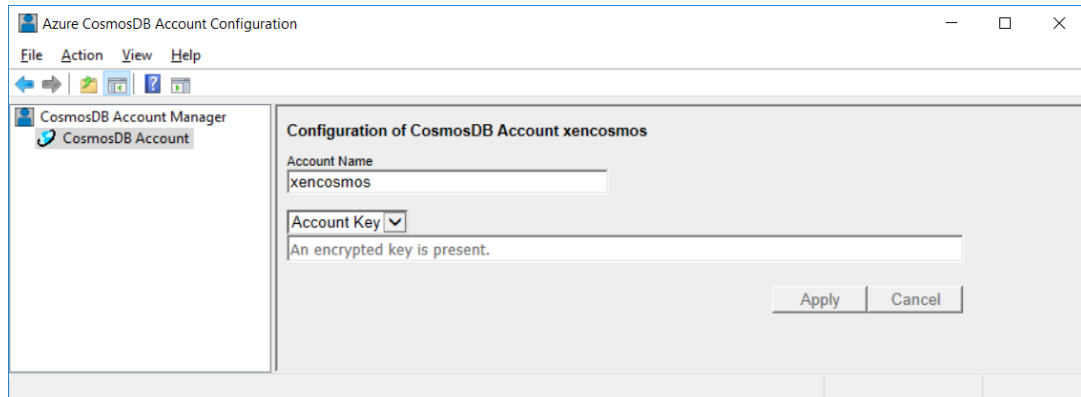
1. Purchase a Cloud File Gateway with Multi-Site Sync subscription. XenData will then provide subscription activation codes and a Cosmos DB account and key combination.
2. Install the Cloud File Gateway and Multi-Site Sync Extensions to Archive Series software.
3. License the Cloud File Gateway and Multi-Site Sync subscriptions using the License Administration Utility to enter the activation codes provided by XenData.
4. Follow the steps outlined in [Configuring Azure Storage Accounts](#), [Configuring Amazon S3 Endpoint](#) or [Configuring a Wasabi S3 Account](#) to 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. This step should be completed before Adding Cosmos DB Account Access.
5. Add Cosmos DB Account Access, as described in the [Adding Cosmos DB Account Access](#) section.

5.5.2 Adding Cosmos DB Account Access

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

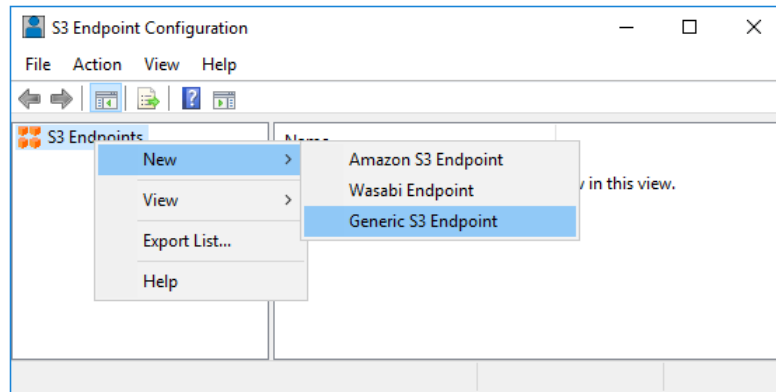
Once you have completed the recommended actions outlined at the top of this page, along with the instructions here, on each of your XenData servers, be they physical or virtual, you will be ready to use the XenData Multi-Site Sync.

5.6 Configuring Generic S3 Endpoints


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

5.6.1 Adding Generic S3 Account Access

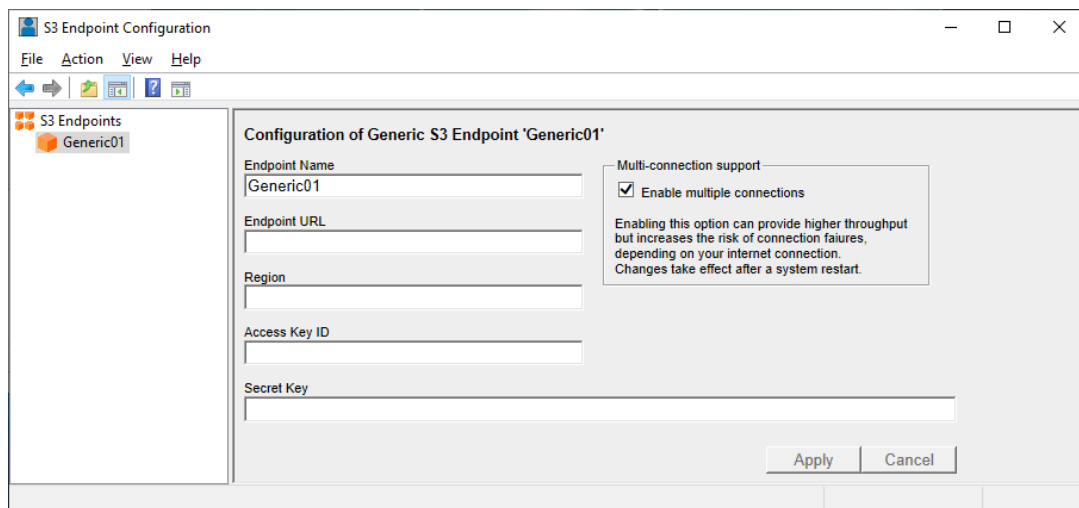
1. Launch the S3 Endpoint Configuration utility as follows:
 1. Click the Windows Start icon.
 2. Open the XenData program group
 3. Click the **S3 Endpoint Configuration** entry in the list.
2. Right-click on 'S3 Endpoints'; select 'New' and then 'Generic Endpoint'.



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



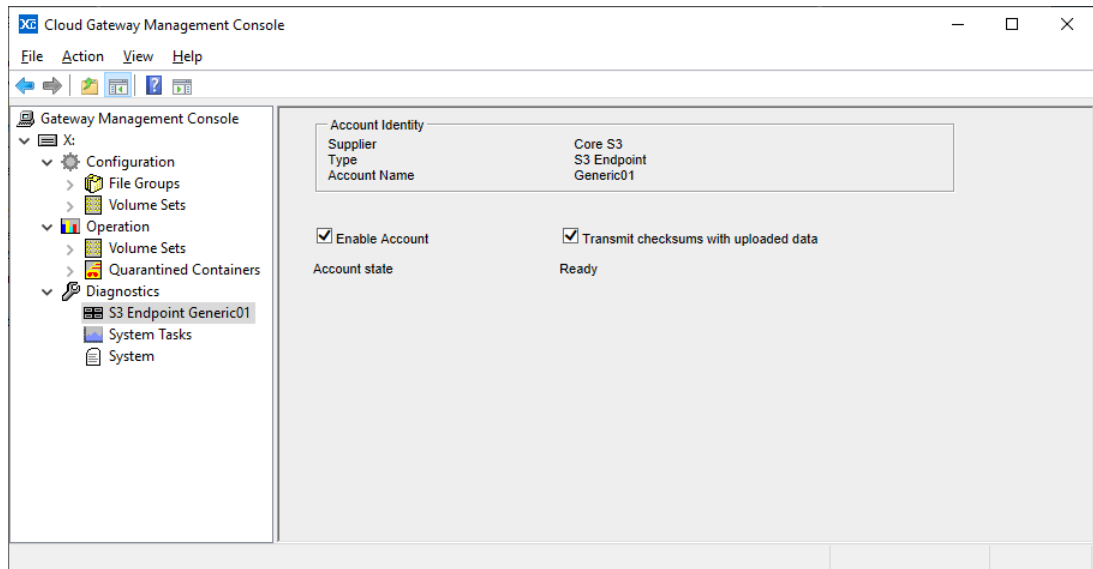
4. Left-click on the endpoint name shown under 'S3 Endpoints', then enter the 'Endpoint URL', 'Region', 'Access Key ID' and 'Secret Key' from your S3 account. Once you've entered all your account information click 'Apply'. You can optionally enable 'Multi-Connection support', this allows the system to 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.

5.6.2 Configuring a Generic S3 Account

1. Expand the **Diagnostics** section in the left pane of the Cloud Gateway 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.

5.7 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 Cloud File Gateway software, an initial Volume Set is automatically created ready for configuration, as illustrated below.



5.7.1 Creating a New Volume Set

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

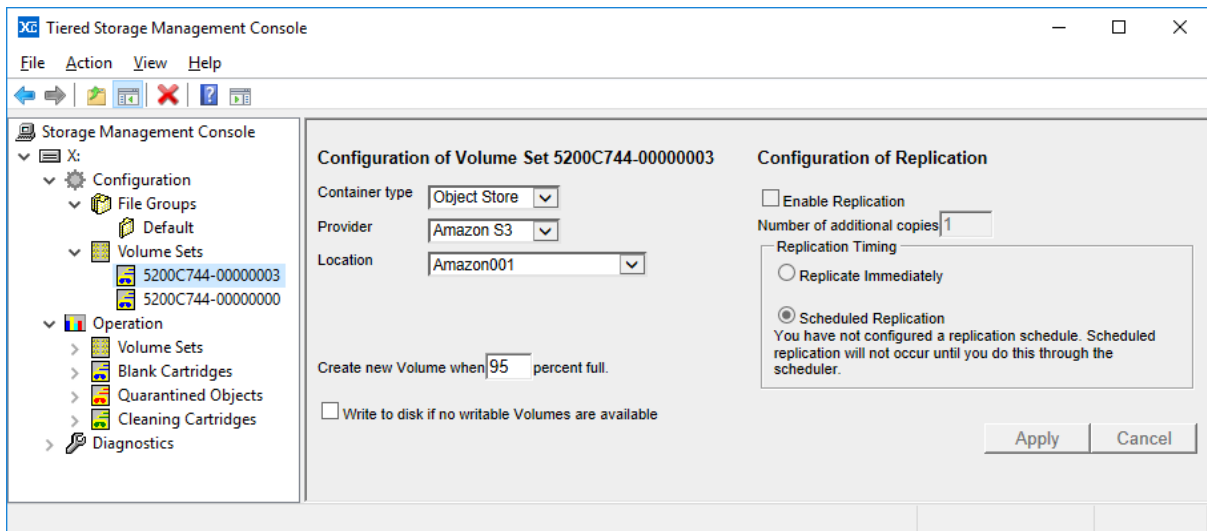
The new Volume Set is now ready to be configured, as described in [Configuring a Volume Set for Object Storage](#) and [Configuring a Volume Set for LTO or ODA](#).

5.7.2 Renaming a Volume Set

1. Expand the **Configuration** section in the left pane of the Cloud Gateway 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**

5.7.3 Configuring a Volume Set for Cloud Storage

1. Expand the **Configuration** section in the left pane of the Cloud Gateway 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](#).

5.7.4 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 Cloud Gateway 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.

5.7.5 Adding a Volume

1. Expand the **Operation** section in the left pane of the Cloud Gateway 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.

5.7.6 Deleting a Volume Set

1. Expand the **Configuration** section in the left pane of the Cloud Gateway 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](#).

5.7.7 Deleting a Volume

1. Expand the **Operation** section in the left pane of the Cloud Gateway 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.

5.7.8 Rebuilding Volume Contents Catalogs

1. Expand the **Operation** section in the left pane of the Cloud Gateway 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](#).

5.7.9 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 Cloud Gateway 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 Cloud Gateway Management Console
2. Expand the **Volume Sets** section
3. Right click on the required Volume Set
4. Click on **Import Folder Structure**

5.7.10 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 Cloud Gateway 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 Cloud Gateway Management Console
2. Expand the **Volume Sets** section
3. Right click on the required Volume Set
4. Click on **Import Data**

5.7.11 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 Cloud Gateway 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

5.7.12 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 Cloud Gateway 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.

5.7.13 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 Cloud Gateway 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.

5.8 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 Cloud File Gateway 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](#).

5.8.1 Creating a New File Group

To create a new File Group:

1. Expand the **Configuration** section in the left pane of the Cloud Gateway 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](#).

5.8.2 Renaming a File Group

To rename a File Group, as displayed in the Cloud Gateway Management Console:

1. Expand the **Configuration** section in the left pane of the Cloud Gateway 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

5.8.3 Changing the Order of File Groups


The order of File Groups in the Cloud Gateway 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 Cloud Gateway 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**

5.8.4 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 Cloud Gateway 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.

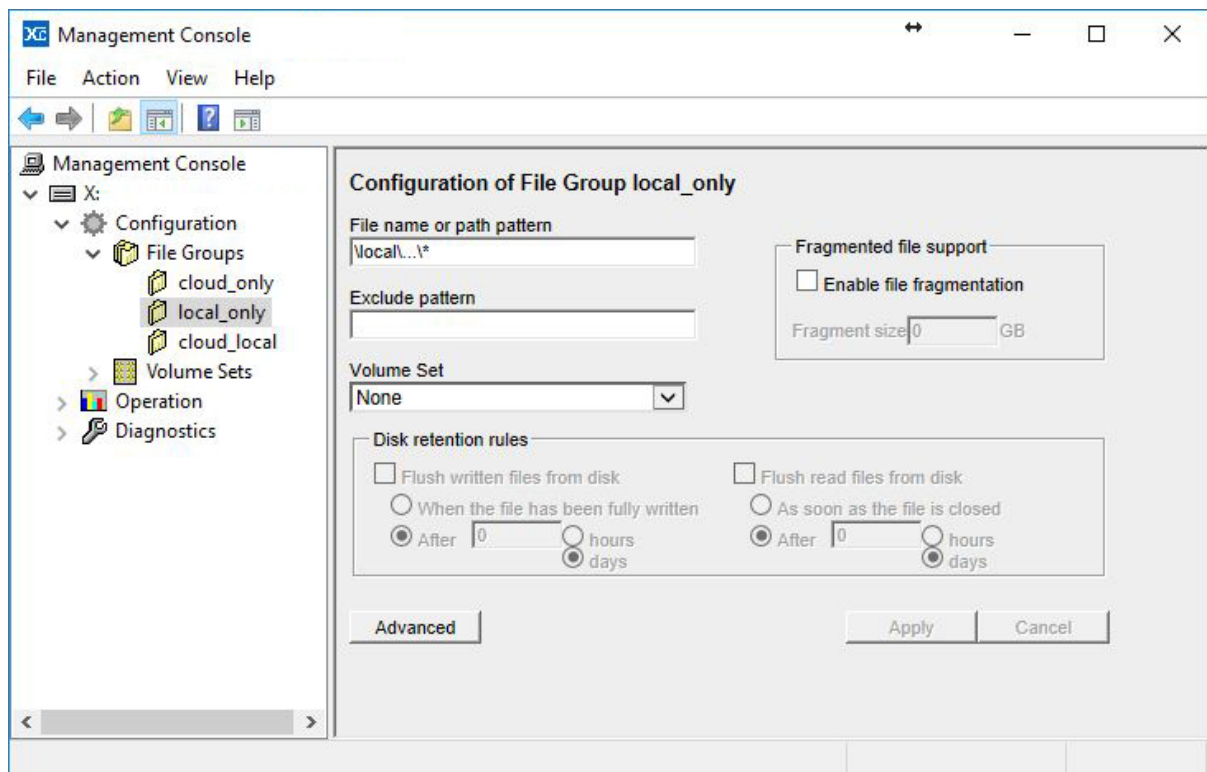
The order of File Groups in the left pane of the console is important and affects how files are allocated to File Groups (see [Changing the Order of File Groups](#)).

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.

5.8.5 Selecting a Volume Set for a File Group

1. Expand the **Configuration** section in the left pane of the Cloud Gateway 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.



5.8.6 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 Object Storage by simply reading the file.

To configure disk retention rules:

1. Expand the **Configuration** section in the left pane of the Cloud Gateway 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, whichever 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.

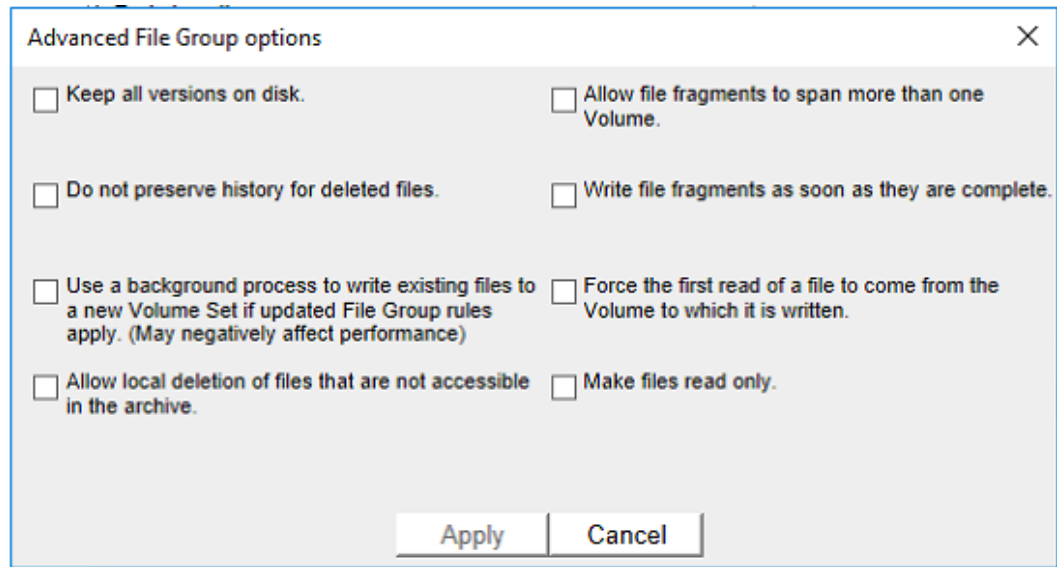
5.8.7 Changing Disk Retention Rules

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.

5.8.8 File Group Advanced Options

To configure advanced File Group options:

1. Expand the **Configuration** section in the left pane of the Cloud Gateway 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 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 Object Storage, even if it is available from cache disk, thereby allowing applications to verify the integrity of data written to Object Storage. **Note** that only the first read is forced to come from the 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.

5.9 S3 Server Interface

5.9.1 Adding the S3 Server Interface

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

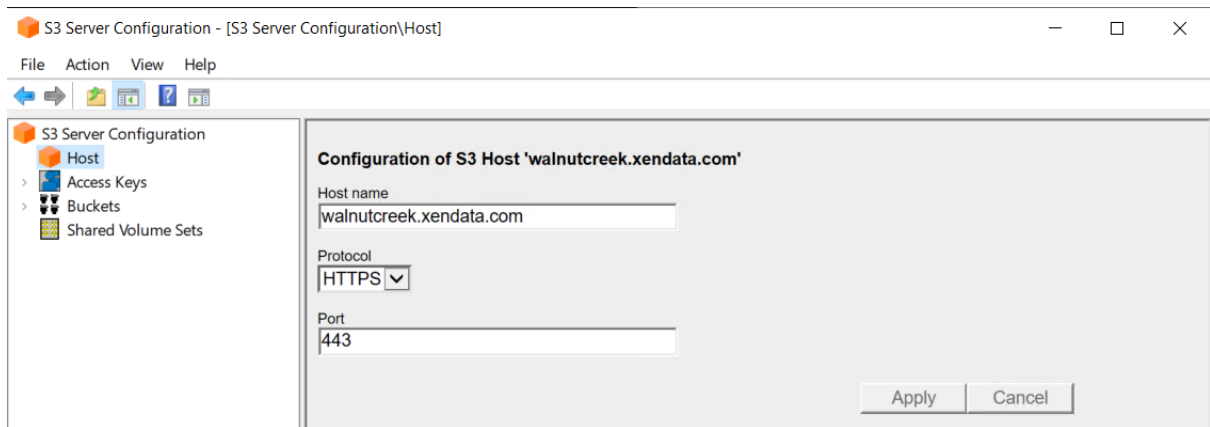
The following are required to add the S3 Server interface to an existing XenData installation.

- ❖ 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.

5.9.2 Configuring the S3 Server Interface

5.9.2.1 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.

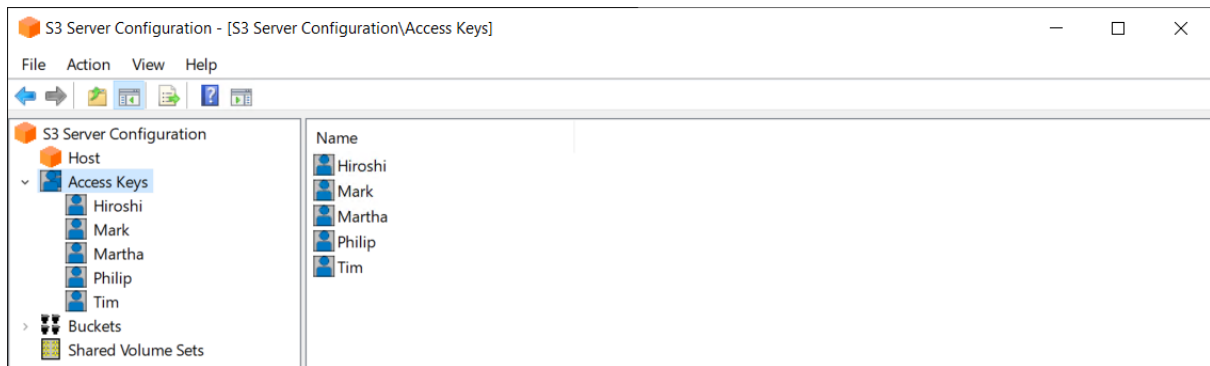


3. After the variables have been entered, left click on 'Apply',.

5.9.2.2 Access Keys

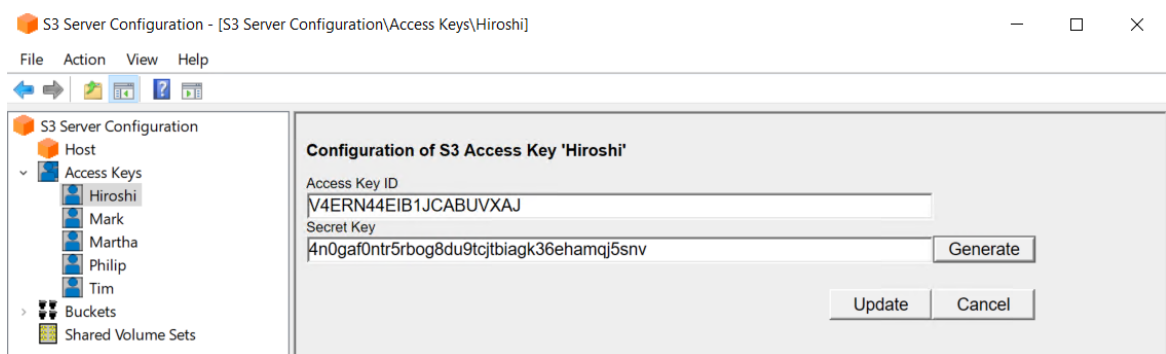
Access Keys are used to allow external applications to authenticate their access to the XenData S3 Server Interface. To generate a new access key:

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. 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.
3. 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'.



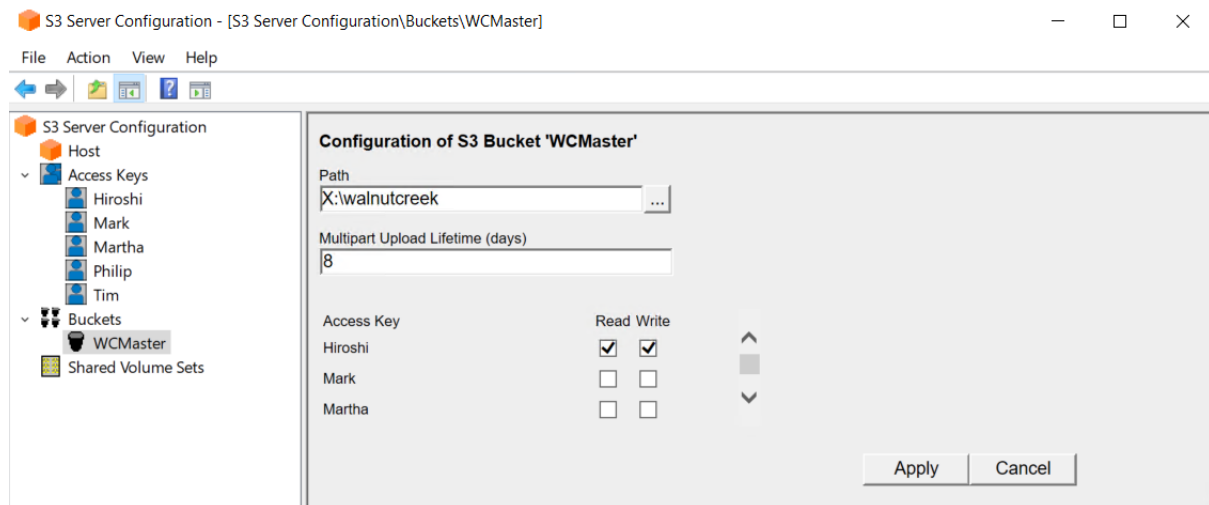
4. Once satisfied with the Access Key and Secret Key, select 'Update'.
5. Repeat the above steps to create additional access key and secret key pairs.

5.9.2.3 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.

6. File Explorer Plug-In

On the computer running Cloud File Gateway 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.

6.1 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 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 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 Cloud File Gateway 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 Cloud Gateway Management Console are not set to flush immediately after a file is closed, this will result in the file being fetched back to disk.

6.2 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 Cloud File Gateway 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 Cloud Gateway 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 Cloud File Gateway 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.

6.3 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.

6.4 Enhanced Properties

6.4.1 Obtaining Enhanced Properties

Enhanced properties are available for the logical drive managed by the Cloud File Gateway software as described below.

To obtain Enhanced Properties:

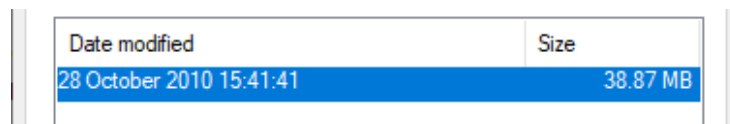
1. Open Windows File Explorer on the computer running the Cloud File Gateway software.
2. Right click on the logical drive letter under XenData control.
3. Select Properties and then select the XenData tab.

6.4.2 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.



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

Change Tier X

Container	Tier location	Available tiers	Rehydrate priority
492a7988-00000002-6137881e	Archive	Not set	Standard

- 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.

Container	Tier location	Available tiers	Rehydrate priority
492a7988-00000002-6137881e	Archive	Cool	Standard
			Standard
			High

- 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.

6.5 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 Cloud File Gateway software.
2. Select Volume View in the left navigational pane.
3. Browse the Volume View.

6.6 History Explorer for Cloud File Gateway

The default behavior of the XenData Cloud File Gateway does not retain old file versions and deleted files on Object Storage. However, it may be configured to do so via a registry setting. In this case, History Explorer may be used to list all available versions of all files, all file instances and their Volume locations, including deleted and renamed files. It also allows the retrieval of old, overwritten or deleted file versions.

To Browse with History Explorer

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

7. 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 Cloud Gateway 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](#).

7.1 Starting Metadata Backup

To start the program:

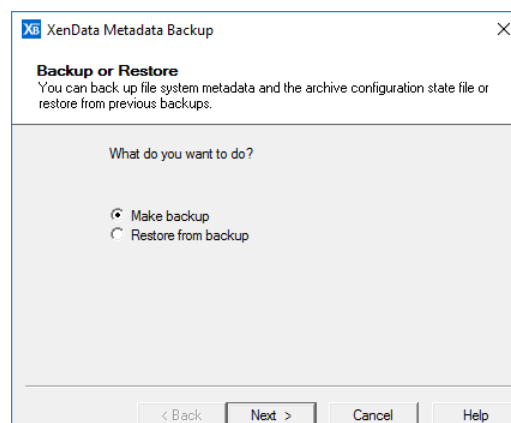
1. Click the Windows Start icon
2. Open the XenData program group
3. Click the **XenData Metadata Backup** entry in the list

7.2 Selecting Backup or Restore

The Metadata Backup program performs two types of operation:

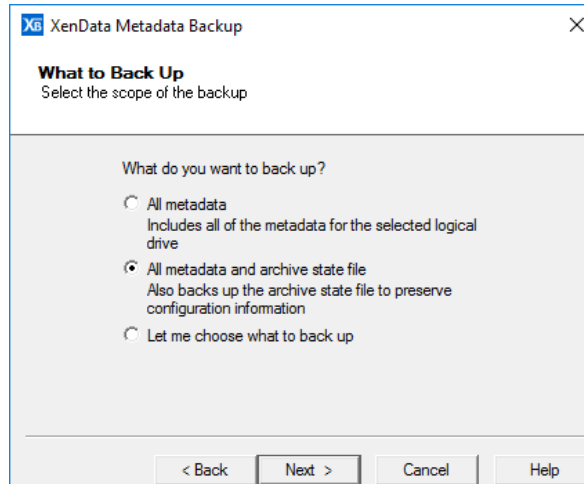
- ❖ Make backup - makes a backup of the metadata in the system in its current state. See 'Making a Predefined backup' or 'Making a Custom Backup' below.
- ❖ Restore from backup – restores metadata from a backup file onto the cache disk volume. See 'Restore from backup' below.

Select the desired option and click Next to continue.



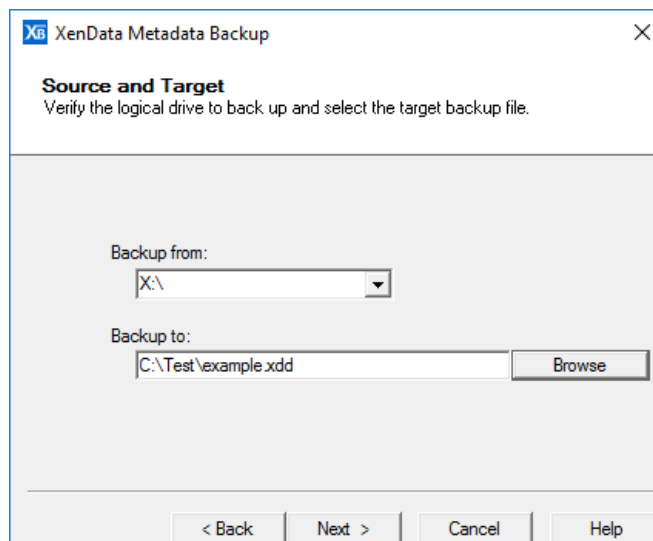
7.3 Making a Predefined Backup

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

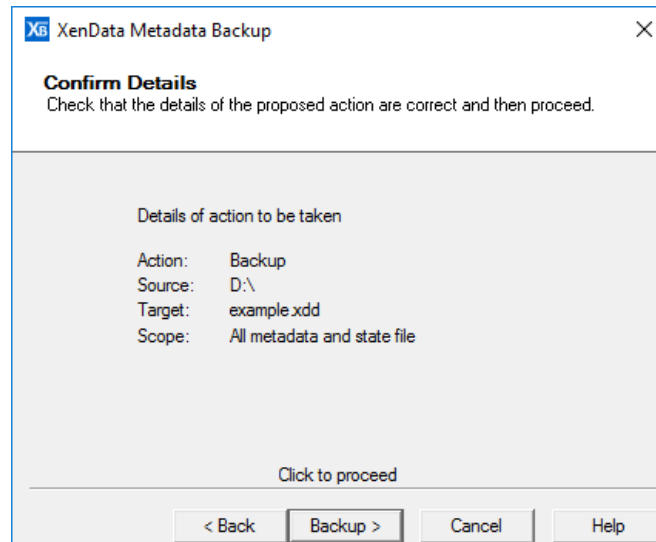


There are two predefined backup types. **All metadata** will back up all the file system metadata, and **All metadata and XenData state file** will also include the XenData state file.

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.



7.4 Making a Custom Backup

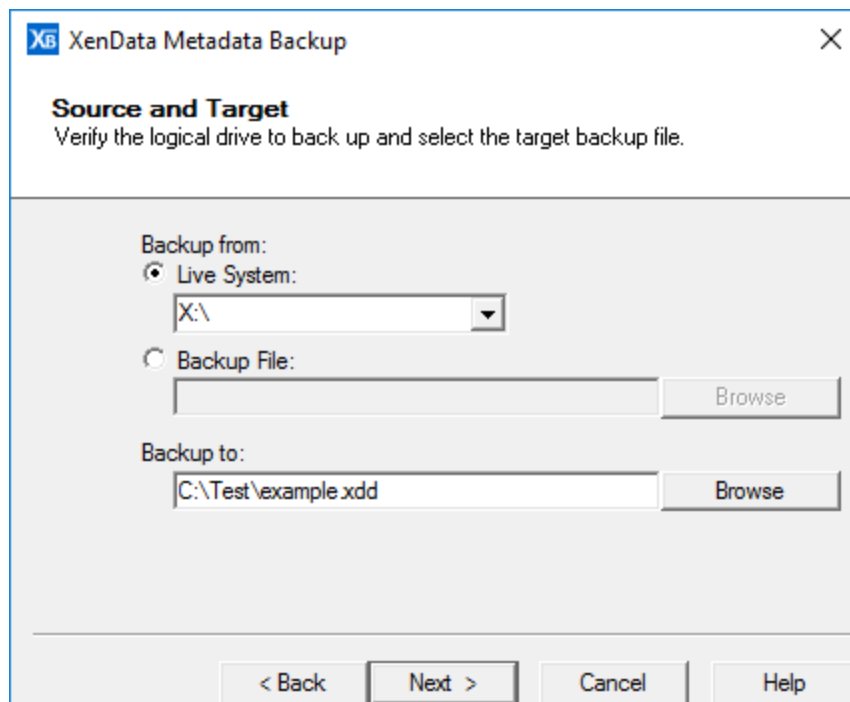
The instructions in this section describe how to perform a partial metadata backup, selecting what is included in the backup. For example, a folder only used for temporary files may be excluded from the backup as the files it contains will not be needed following a system restore. It is also possible to create a sub-backup. This refers to creating a new backup file from an existing backup where the new backup contains only selected folders from the original backup file.

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




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



7.5 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.

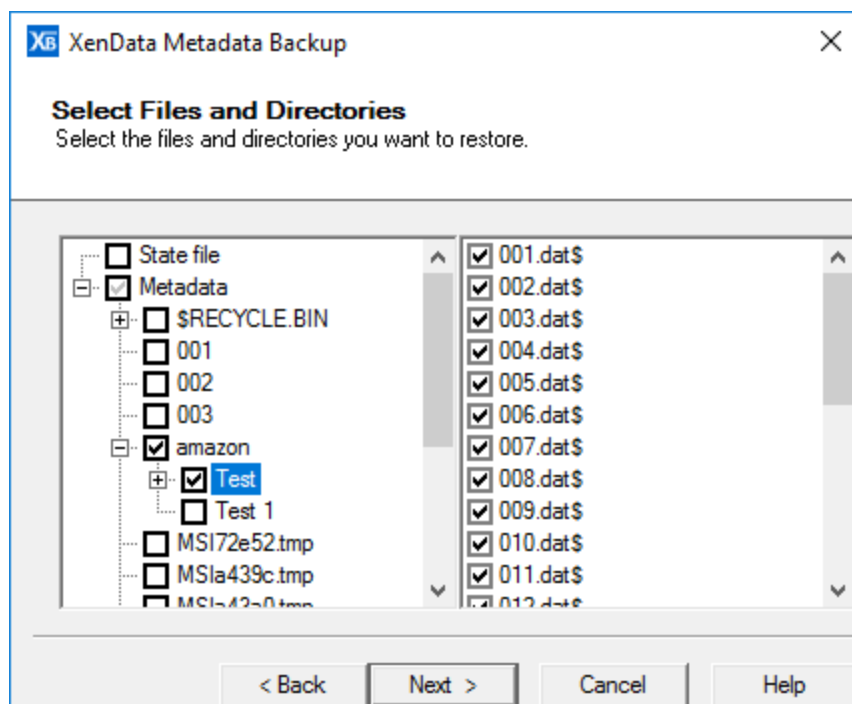


There are three restore options:

- ❖ 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:

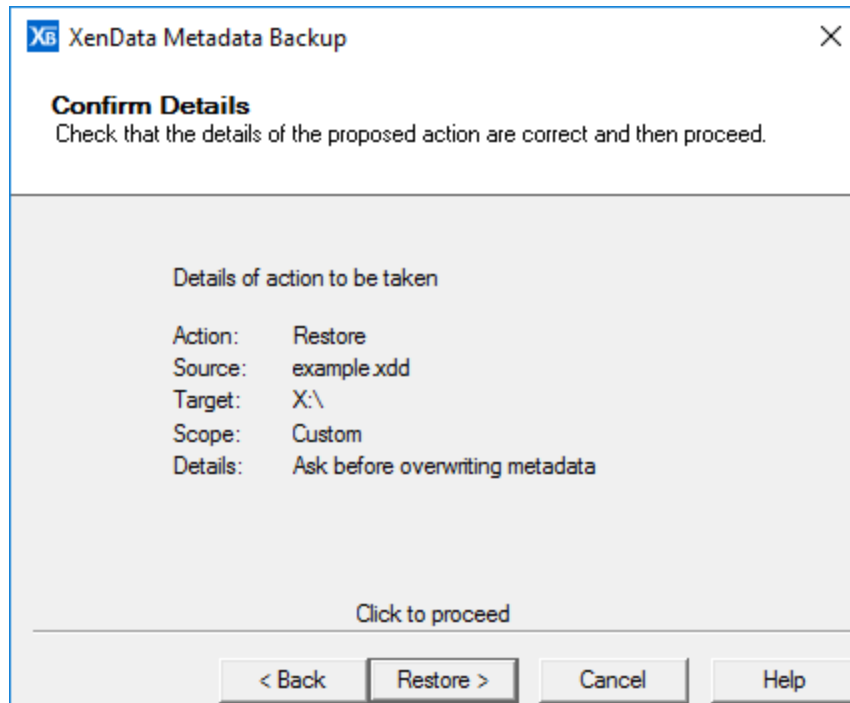
1. Specify the input backup file to restore from, or verify that the correct file name has been determined automatically.
2. Verify the logical drive letter to restore to.
3. Select the desired restore option.
4. Click Next to continue



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.

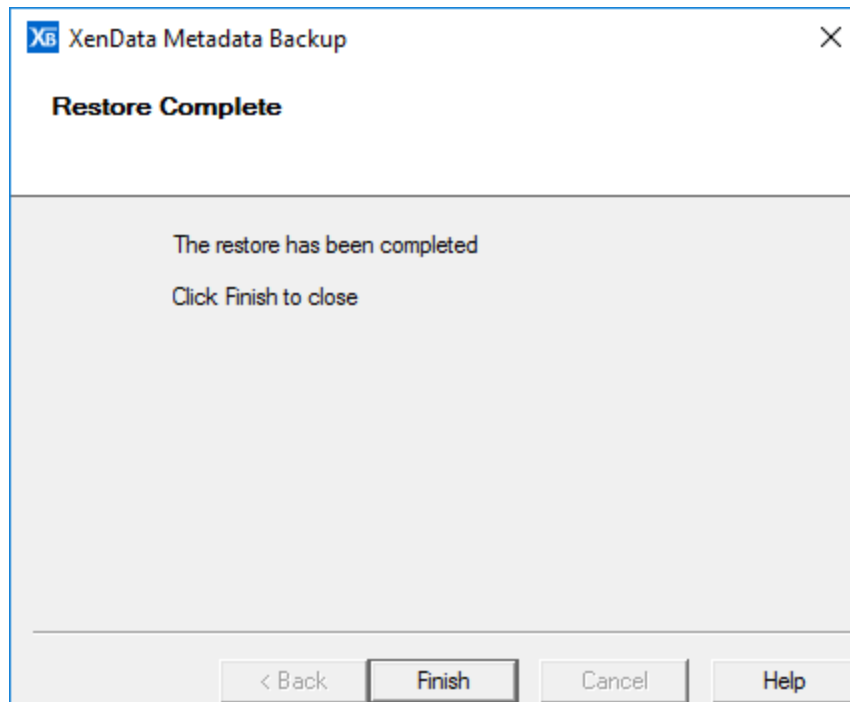
A progress dialog box will appear so that you can check the status of the restore operation. If the option to **Ask before overwrite** was selected during restore configuration, dialog boxes similar to the one shown below might appear, asking if existing metadata should be overwritten, and giving a category 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.



8. 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 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.

8.1 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

8.2 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.



8.3 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.

8.4 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.

8.5 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.

8.6 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:** Checked.

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.

8.7 Scheduling Deferred Write

DeferredWrite

Recurrence: None Daily Weekly Monthly

Hourly: [dropdown]

Start: [date: 2020-01-03] [time: 21:15] Expire [date: 2021-12-31] [time: 21:15]

Task Name: [text: After Hours Update] Stop task if it runs longer than [dropdown]

Number of drives to use for deferred writes: [dropdown: 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.

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

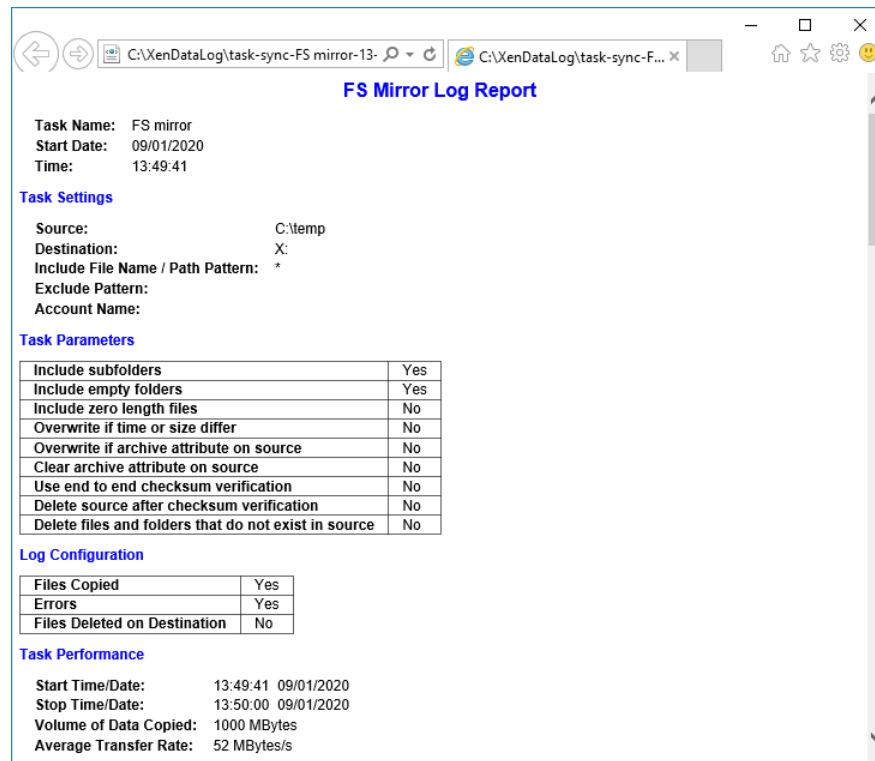
8.8 Scheduling File System Mirror

The screenshot shows the 'FS Mirror' configuration window with the following details:

- Recurrence:** Radio buttons for None, Hourly, Daily (selected), Weekly, and Monthly. A dropdown menu is next to the Daily option.
- Start:** Date: 2020-01-03, Time: 15:45. An 'Expire' checkbox is present and unchecked.
- Task Name:** Text field containing 'FS Mirror'. A 'Stop task if it runs longer than' checkbox is present and unchecked.
- Logging and File Handling:**
 - Use Log File
 - Log Errors
 - Include subfolders
 - Log all copied files
 - Log deletions on destination
 - Include empty folders
 - Include zero length files
 - Overwrite if size or time differ
 - Overwrite if source has archive attribute set
 - Clear archive attribute on source
 - Use end-to-end checksum verification
 - Delete source after checksum verification
 - Delete files and folders that do not exist in source
- Source Folder:** Text field containing 'C:\Test' with a folder selection icon.
- Destination Folder:** Text field containing 'X:\FS Mirror' with a folder selection icon.
- Include file name or file path pattern:** Text field containing '*'.
- Exclude Pattern:** Empty text field.
- User Account:** Empty text field.
- Password:** Empty text field.
- Buttons:** 'Check', 'Save', and 'Cancel' buttons.

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 of 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.

8.9 Scheduling File System Mirror Reporting Run

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.

The screenshot shows a web browser window displaying an 'FS Mirror Log Report'. The report details the execution of a task named 'FS mirror' on 09/01/2020 at 13:49:41. It lists task settings such as source (C:\temp) and destination (X:), and task parameters including options for subfolders, empty folders, and checksum verification. A 'Log Configuration' table shows that files copied, errors, and files deleted on destination are all recorded. Finally, a 'Task Performance' section provides timing and volume data.

FS Mirror Log Report

Task Name: FS mirror
 Start Date: 09/01/2020
 Time: 13:49:41

Task Settings

Source: C:\temp
 Destination: X:
 Include File Name / Path Pattern: *
 Exclude Pattern:
 Account Name:

Task Parameters

Include subfolders	Yes
Include empty folders	Yes
Include zero length files	No
Overwrite if time or size differ	No
Overwrite if archive attribute on source	No
Clear archive attribute on source	No
Use end to end checksum verification	No
Delete source after checksum verification	No
Delete files and folders that do not exist in source	No

Log Configuration

Files Copied	Yes
Errors	Yes
Files Deleted on Destination	No

Task Performance

Start Time/Date: 13:49:41 09/01/2020
 Stop Time/Date: 13:50:00 09/01/2020
 Volume of Data Copied: 1000 MBytes
 Average Transfer Rate: 52 MBytes/s

- ❖ 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.

9. Reports

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

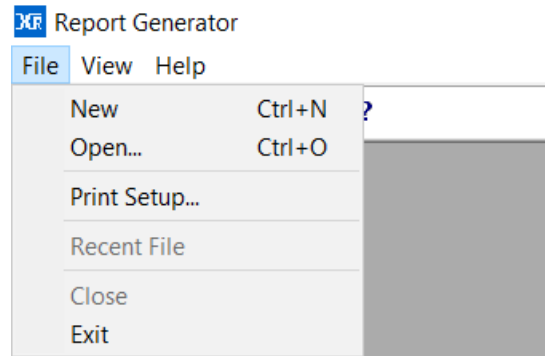
9.1 Starting the Report Generator

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


9.2 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.

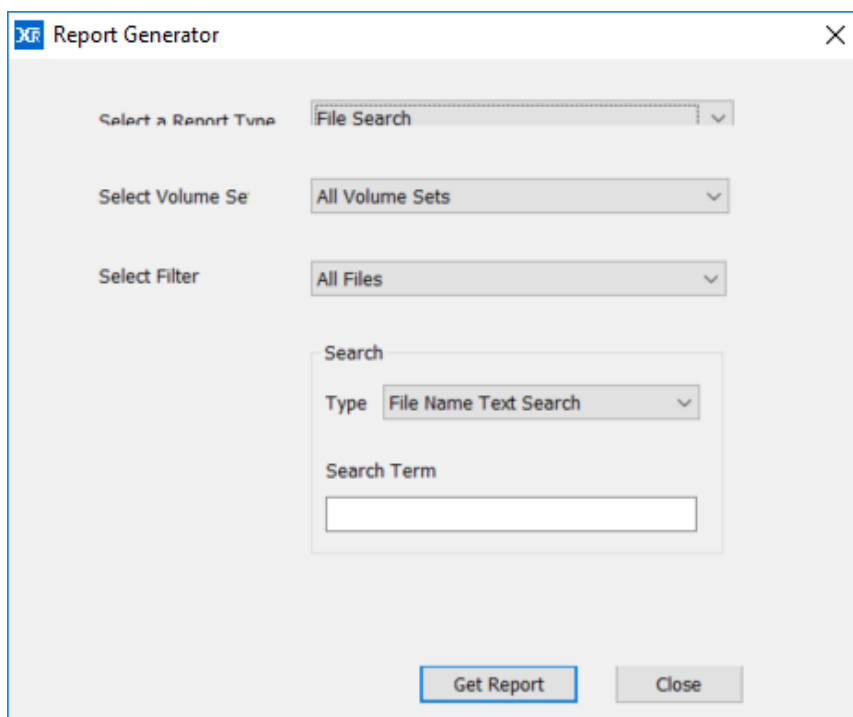


Then browse to the location of the saved report, then select the required XRG file and open it.

9.3 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.

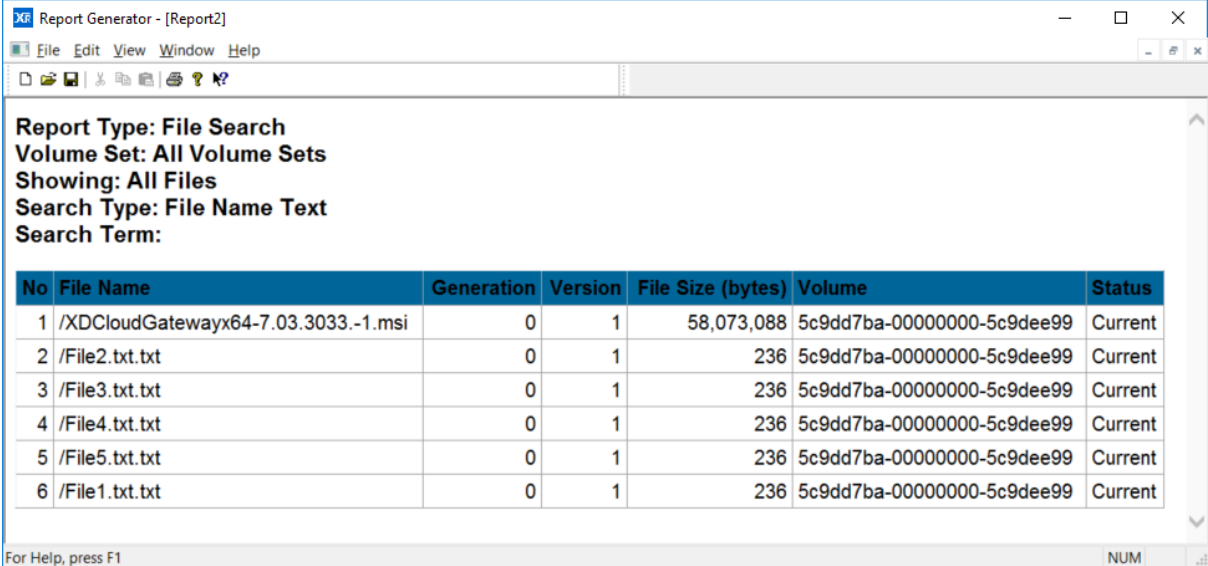
It is possible to search using a simple text search or using a Regular Expression. When **File Name Text Search** is chosen, the search option supports wild cards.

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.

9.3.1 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 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 Volume followed by the total size of the file in bytes.

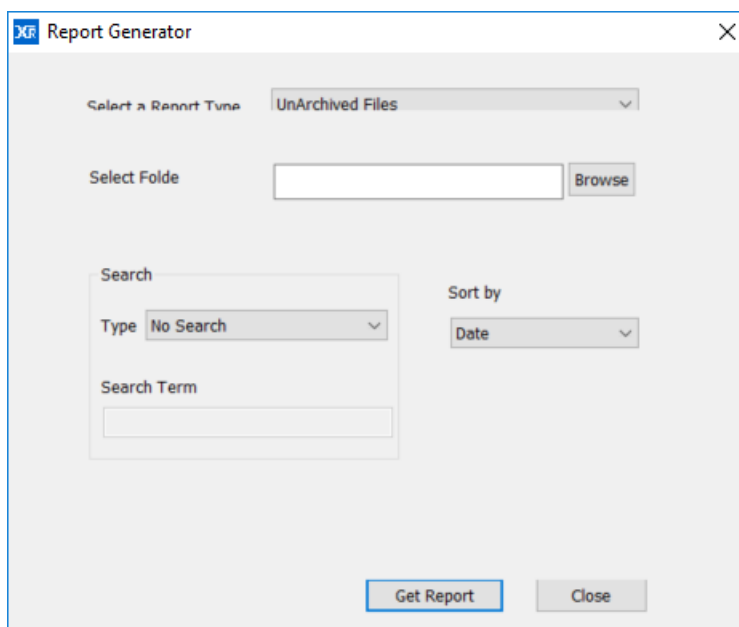
- Volume - this is the Volume that contains the file.
- Type - The status of the file is displayed as one of the following:
 - Current - this is the most recent version of the file, accessible through the archive drive letter.
 - Renamed - the file has been renamed and is now accessible under a different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

9.4 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**.

9.4.1 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 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.

- 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.

9.5 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.

9.5.1 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

10. Alert Module

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

10.1 About the Event Monitor

The Event Monitor runs on the same server as the Cloud File Gateway 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](#).

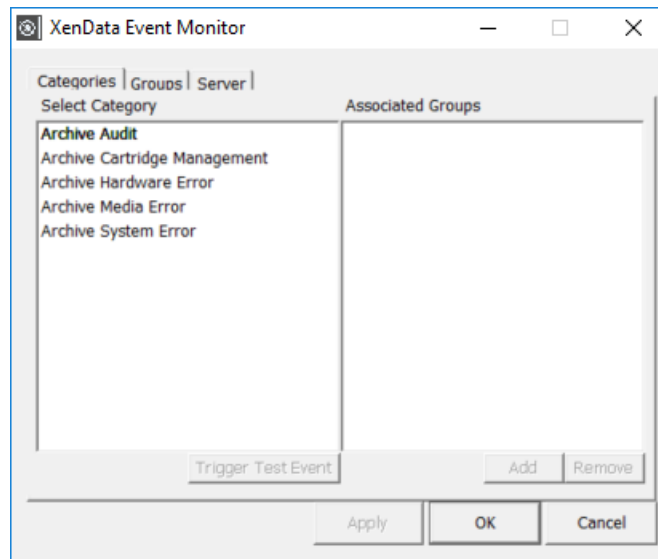
Set up of the Event Monitor is described in [Configuring the Event Monitor](#).

10.2 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.

10.3 About Event Categories

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

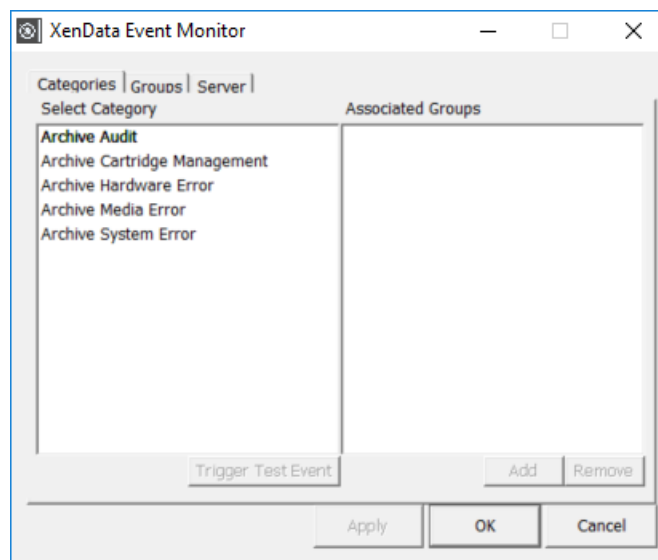
- ❖ **Archive Audit:** This category of event messages describes the successful completion of routine operations.
- ❖ **Archive Media Management:** This category of event messages may require routine action from the gateway operator.

- ❖ 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 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](#).

10.4 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**

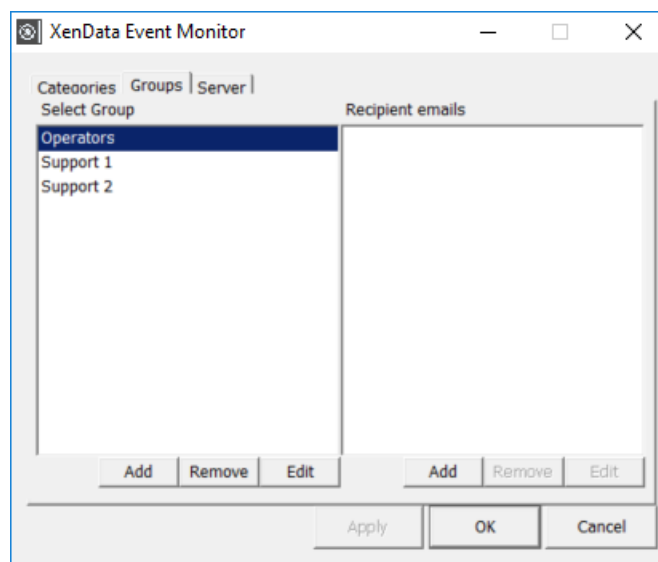
Repeat this mapping for each event category, as required and then click **Apply**.

10.5 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](#).

10.6 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**

10.7 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](#).

10.8 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' text boxes, 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**

10.9 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).

10.10 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 Cloud File Gateway software or a connected Windows client. The On-Screen Messaging and Event Monitor are installed automatically on the machine running the Cloud File Gateway 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.

10.11 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.

11. Diagnostics & Maintenance

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


11.1 Windows Event Log

Whenever the Cloud File Gateway 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 Cloud File Gateway software successfully completed an inventory of the 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.

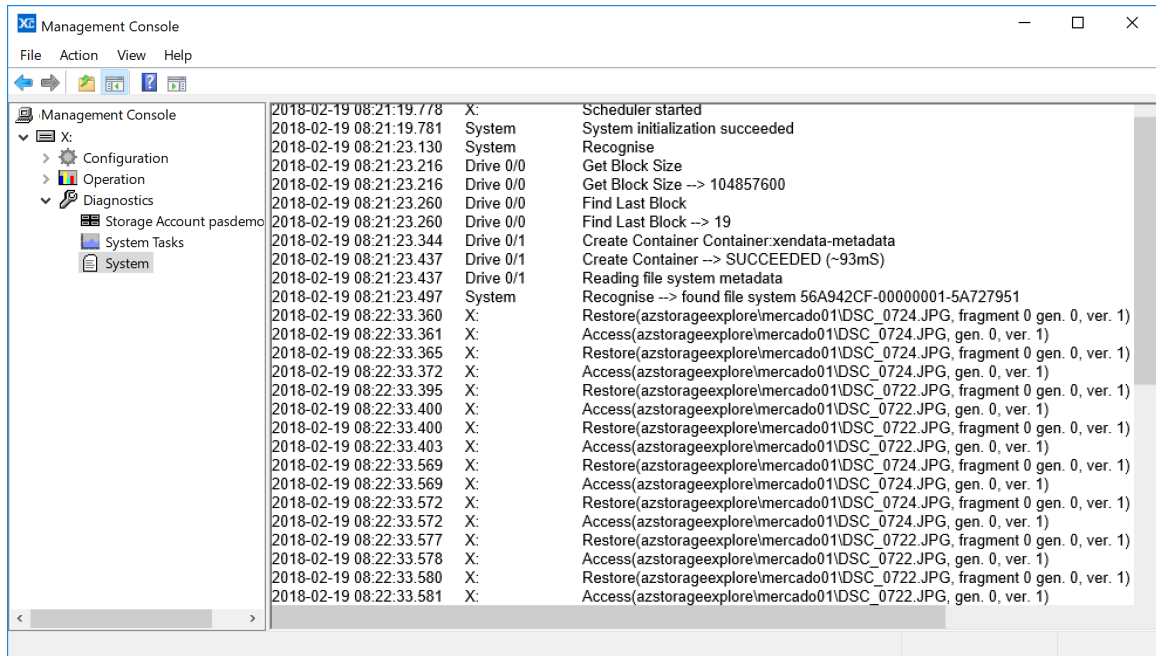


11.2 System Trace Log

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 Object Storage.

To Open the Trace Log

1. Open the Cloud Gateway 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 Cloud File Gateway 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.

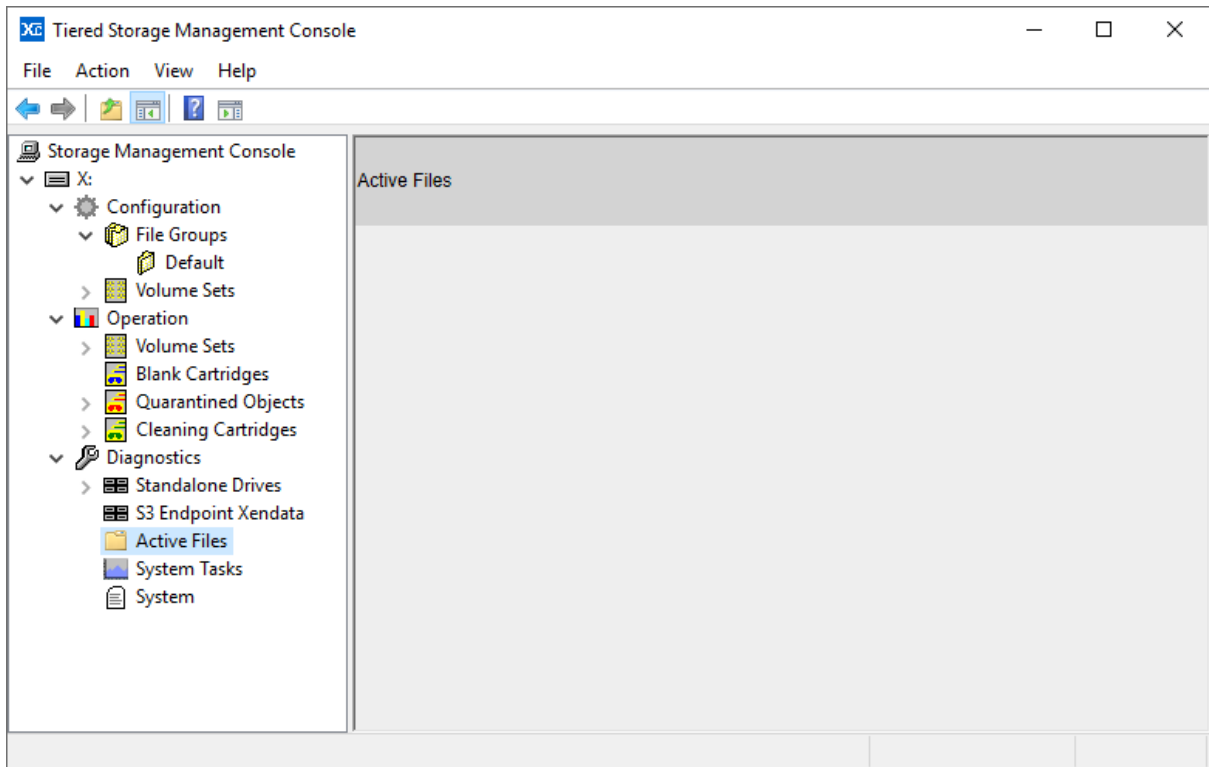
11.3 Active Files Display

The Active Files display lists files that are being actively archived and restored. It shows their progress as they are cached on the archive system and transferred between the archive and the application that initiated the archive or restore operation.

To Open the Active Files Display

1. Open the Cloud Gateway Management Console.
2. Navigate to the Diagnostics section.

- Click on the Active Files folder icon to open the Active Files display in the right pane of the window.



Interpreting the Display

The display uses four different colored progress bars, two for archive operations and two for restores.


For restore operations:

- ❖ yellow displays progress as a file is restored from the archive storage to cache,
- ❖ blue displays the portions of a file that have been consumed by the application.

This example shows a file that is queued for restore but no data has yet been transferred.




The next example below shows the file being restored to cache and simultaneously transferred to an application.



For archive operations:

- ❖ light green shows progress as a file is written to cache,
- ❖ dark green shows progress as a file is written to the archive storage.

The example below shows two files that have been written to cache. The first file is being written to archive storage.



Operations Not Displayed

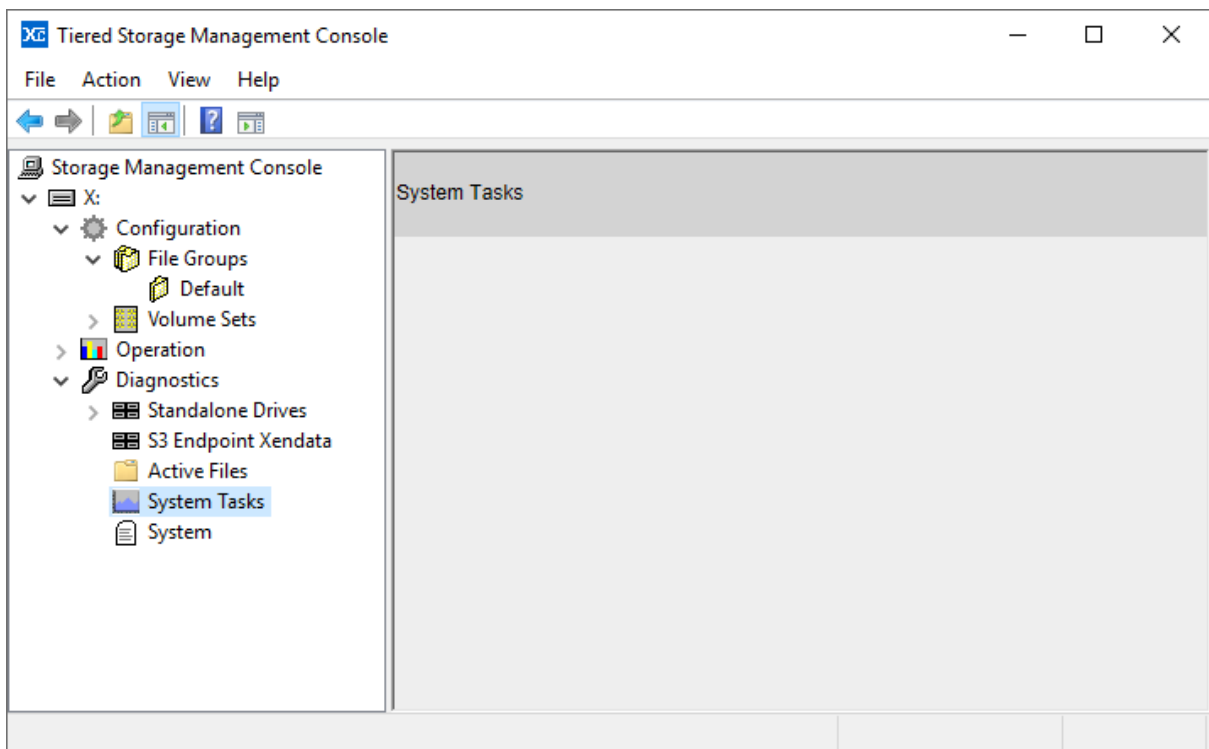
The Active Files display does not show file rename, delete or repack operations. Repack operations are shown in the [System Tasks](#) display.

11.4 System Tasks

The System Tasks display shows progress of a operation being completed on the XenData system. It shows the current progress whilst the operation is running and will disappear after the task is completed.

To Open the System Task Display

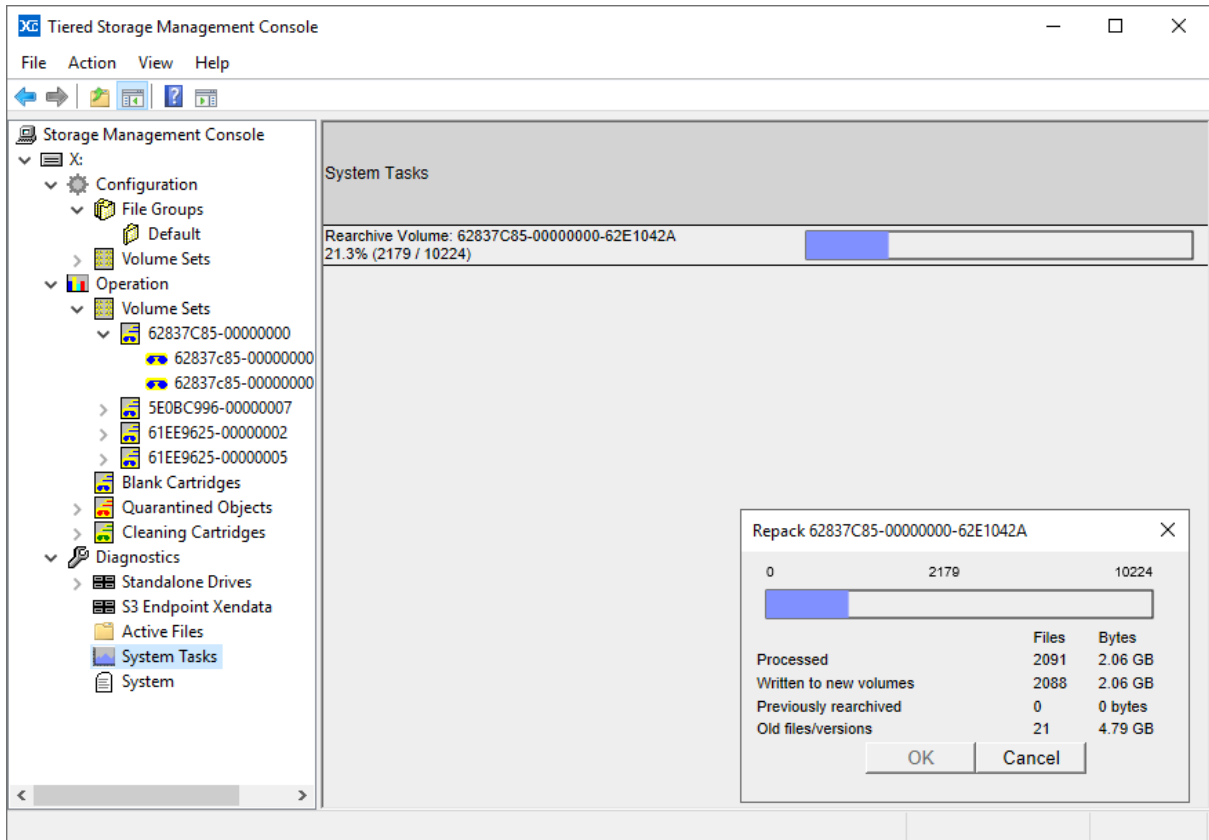
1. Open the Cloud Gateway Management Console.
2. Navigate to the Diagnostics section.
3. Click on the System Tasks icon to open the System Tasks display in the right pane of the window.



The Systems Tasks display shows progress on long operations such as:

- ❖ [Rebuild Catalog](#)
- ❖ [Import Folder Structure](#)
- ❖ [Import Data](#)
- ❖ [Volume Statistics](#)

The example below shows a repack operation being completed in System Tasks



12. 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 Cloud File Gateway software. There are three utilities that may be installed:

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

12.1 Installation Prerequisites

Please ensure that the following prerequisites are met:

- ❖ Apply registry setting to the computer running XenData Cloud File Gateway as described in XenData Technical Note XTN1803. This is required when the Cloud File Gateway software is running on Windows Server 2016 or Windows Server 2012 R2.
- ❖ The client computer must be connected to the same network as the computer running the Cloud File Gateway software. They may be connected within a domain or a workgroup.
- ❖ The XenData share must be mapped to a drive letter on the client computer.

12.2 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'.

12.3 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.

12.4 File Explorer Plug-In

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](#)

12.5 Trace File Viewer

Whenever XenData Cloud File Gateway 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.

12.6 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.

13. Glossary

Activation Code An Activation Code is required to run the Cloud File Gateway software and enables the chosen configuration. Separate Activation Codes are required to enable Multi-Site Sync, FS Mirror and Alert Module functionality. The Cloud File Gateway 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 Cloud File Gateway 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 Cloud File Gateway 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.

Cloud Gateway 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 Cloud File Gateway 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 Cloud File Gateway 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 Cloud File Gateway 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 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 Cloud File Gateway software can be configured to automatically flush files from the disk cache once they are securely stored on 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 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 Cloud File Gateway 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 Cloud Gateway 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 or, two or more Archive Series instances with Cloud File Gateway Extension, Sync Service and S3 Instance enabled, sharing volume sets. The Multi-Site Sync 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 Cloud File Gateway 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 Cloud File Gateway 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 Cloud Gateway 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 Cloud File Gateway 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.