Quest

Quest® NetVault® Plug-in *for FileSystem* 14.0.1

User's Guide

© 2025 Quest Software Inc.

ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept.
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our website (https://www.quest.com) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at https://www.quest.com/legal.

Trademarks

Quest Software, Quest, the Quest logo, QoreStor, and NetVault are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit https://www.quest.com/legal/trademark-information.aspx. All other trademarks and registered trademarks are property of their respective owners.

Legend

WARNING: A WARNING icon indicates a potential for property damage, personal injury, or death.

CAUTION: A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

IMPORTANT NOTE, **NOTE**, **TIP**, **MOBILE**, or **VIDEO:** An information icon indicates supporting information.

NetVault Plug-in *for FileSystem* User's Guide
Updated - August 2025
Software Version - 14.0.1
NVG-101-140-EN-01

Contents

Introduction 5

About Quest® NetVault® 5

Key benefits 5

Feature summary 5

Target audience 6

Recommended additional reading 6

Installing the plug-in 7

Installing the plug-in 7

Installing the plug-in using the push installation method (Windows) 7

Installing the plug-in using the configuration wizard (Linux/UNIX) 8

Installing the plug-in from the Manage Clients page 8

Removing the plug-in 9

Using the correct plug-in binary on Solaris 9

Configuring the plug-in 10

Configuring default settings 10

Exclusion lists 12

About exclusion lists 13

Exclusion list considerations 14

Backing up data 15

Backing up with Plug-in *for FileSystem* 15

About backup methods 15

About backup types 16

About System State Backups 16

About multistream backups 17

Backup and recovery strategy 17

Performing Full Backups 18

Selection tree icons: Linux and UNIX 24

Selection tree icons: Windows 25

Inclusion lists 26

About inclusion lists 26

Inclusion list considerations 26

Adding an inclusion list 27

Removing an inclusion list 27

Additional notes 27

Performing Incremental and Differential Backups 30

Stopping a job 32

Restarting a job 32

Performing Network Share Backups 33

Prerequisites 33

Configuring a network share 34

Configuring a login account for network share . . . 34
Backing up network share data . . . 34
Performing System State Backups . . . 35
System State backups: Additional notes . . . 35

Restoring data . . . 37
Restoring data using the plug-in . . . 37
Searching for files in savesets . . . 39
Restoring an item using the Quick Save feature . . . 40
Viewing media list . . . 41
Renaming and relocating data during restore . . . 41
Restoring data to an alternate client . . . 42
Prerequisites . . . 42
Restore procedure . . . 42
Restoring System State Backups . . . 42
About System State restore . . . 42
System State restore considerations . . . 43
Supported restore methods for Active Directory . . . 43
Supported restore methods for SYSVOL Directory . . . 44
Prerequisites . . . 44
Restoring System State data . . . 45
Restoring cluster database on Windows 2008 . . . 46
Additional considerations for domain controllers . . . 46

Troubleshooting . . . 47
Common errors . . . 47

About us . . . 49
Technical support resources . . . 49

1

Introduction

- About Quest® NetVault®
- Target audience
- Recommended additional reading

About Quest® NetVault®

Quest NetVault (NetVault) increases confidence in the recoverability of critical file system data and eliminates the need for writing complex backup and recovery scripts. Through an intuitive user interface and automated workflow process, the plug-in offers a centralized console to set up, configure, and create flexible backup policies that account for multiple recovery scenarios. Support for Full, Incremental, and Differential Backups allows you to choose a preferred backup strategy. The plug-in minimizes downtime by allowing you to restore full volumes, individual partitions, or individual directories and files quickly and reliably with minimal interaction. Through automatic integration with a wide range of backup devices, you can be confident that your data is protected and safely stored off-site to meet your disaster recovery and business continuity goals.

NOTE: The NetVault is packaged with the NetVault software and automatically installed on the respective machines when you install the NetVault Server or Client software.

Key benefits

- **Increases confidence and reduces risks for critical data:** The NetVault provides total protection to your NetVault Server and Clients without requiring you to write complex backup scripts. The plug-in allows Full, Incremental, and Differential Backups of file system data, and additionally supports snapshot-based point-in-time consistent backups and System State Backups on Windows-based systems. By relying on the NetVault to implement your backup policies, you can be confident that your data is protected, and you are able to recover exactly what you need in the unfortunate event of a failure.
- **Speeds up restores and reduces downtime:** With the NetVault, backup administrators are no longer required to write scripts or run restore commands to recover lost data. The point-and-click feature reduces manual steps and errors during data recovery, and flexible recovery options allow you to recover full volumes or individual files to the same or alternate location.
- **Ensure business continuity with automatic integration of backup devices:** The NetVault supports a wide-range of disk-based devices, tape-based devices, and physical libraries and standalone drives. Off-site storage to tapes and disks gives you the confidence that your data is protected for disaster recovery purposes. With the NetVault, the administrators need not be present 24x7. Even less-experienced administrators can initiate restores with the certainty that they are performed correctly and quickly, thus reducing downtime and improving business continuity.

Feature summary

- Protects server and client systems
- Supports standard and snapshot-based backup methods

- Allows Full, Incremental, and Differential Backups while data is online and accessible
- Supports multistream backup operations
- Restores volumes, partitions, and individual directories and files
- Allows renaming of files and directories during restores
- Restores data to alternate location
- Restores data to alternate NetVault Client
- Allows System State backups and restores on Windows
- Supports full and incremental backups to QoreStor devices

Target audience

This guide is intended for backup administrators and other technical personnel who are responsible for designing and implementing a backup strategy for the organization. A good understanding of the operating systems under which the NetVault Server and Clients are running is assumed.

Recommended additional reading

- *Quest NetVault Installation Guide:* This guide provides information about installing the NetVault Server and Client software.
- *Quest NetVault Administrator's Guide:* This guide provides information about configuring and using NetVault to protect your data. It provides comprehensive information about all NetVault features and functionality.
- *Quest NetVault CLI Reference Guide:* This guide provides information about using the NetVault command-line utilities.

You can download these guides from https://support.quest.com/technical-documents.

2

Installing the plug-in

- Installing the plug-in
- Removing the plug-in
- Using the correct plug-in binary on Solaris

Installing the plug-in

The NetVault is automatically installed when you install or upgrade the NetVault Server or Client software. The procedures described in this section can be used to reinstall the plug-in if it was removed from a NetVault machine or upgrade the plug-in when a standalone patch is distributed for it.

This section includes the following topics:

- Installing the plug-in using the push installation method (Windows)
- Installing the plug-in using the configuration wizard (Linux/UNIX)
- Installing the plug-in from the Manage Clients page

Installing the plug-in using the push installation method (Windows)

On Windows-based machines, you can use the push installation method to install the plug-ins on several machines at the same time. You can perform push installations from the NetVault WebUI.

Before you start the push installation procedure, verify that the following requirements are met:

- **Copy the packages to a shared location:** Copy the client and plug-in packages to a shared location. Only CIFS shares are currently supported as package stores. The path must be accessible to the NetVault Server and all target machines where you want to install the packages.

  Make sure that you use the original names of the installation packages. Renamed packages cannot be selected for push installations.

- **Configure a package store in NetVault:** After copying the installation packages, configure the shared location details in NetVault. For more information, see the *Quest NetVault Administrator's Guide*.

To install the plug-in on Windows-based clients:

1. In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Install Software/Add Clients**.
2. On the **Select Software/Add Clients** page, select **Install NetVault software on remote machines**.
3. In the **Package Store** list, select the repository that contains the installation packages that you want to deploy.
4. To add plug-in packages, click **Add NetVault plug-in package**.

5 In the **Select Packages for Deployment** dialog box, select the check boxes corresponding to the **".npk"** binary files that you want to use, and click **OK**.

6 Click **Next**.

7 On the **Machines to Have NetVault Software Installed** page, click **Choose Machines**, and select **From available machines**.

8 On the **NetVault Machine Details** tab, select the client that you want to add, and click **OK**.

9 To add more machines, repeat Step 7 and Step 8.

10 To submit the task, click **Install Software/Add Clients**.

You can monitor the progress and status of the task from the **Deployment Task Status** page. For more information, see the *Quest NetVault Administrator's Guide*.

Installing the plug-in using the configuration wizard (Linux/UNIX)

On Linux-based machines, you can use the configuration wizard to install the plug-in on multiple clients at the same time.

NOTE: When you use this procedure, make sure that the plug-in binary file is compatible with the client OS and platform.

To install the plug-in on Linux- and UNIX-based clients:

1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Install Plugins**.

2 In the **NetVault Clients** table, select the clients on which you want to install the plug-in.

3 Click **Choose Plug-in File**, and in the browse window, navigate to the location of the **".npk"** installation file for the plug-in (on the installation CD or the directory to which the file was downloaded from the website).

4 Select the platform-specific binary file for the plug-in.

The binary file for Linux and UNIX is named **"nvf-x-x-x-x.npk"** (where **x-x-x-x** represents the version, build, and platform numbers).

5 Click **Next** to begin installation

After the plug-in is installed successfully, a message is displayed.

Installing the plug-in from the Manage Clients page

From the **Manage Clients** page, you can install a plug-in on a single client.

To install the plug-in from the Manage Clients page:

1 In the Navigation pane, click **Manage Clients**.

2 In the **NetVault Clients** table, select the client, and click **Manage**.

3 At the lower-right corner of the Installed Plug-ins table, click the Install Plugin button (+).

4 Click **Choose Plug-in File**, and in the browse window, navigate to the location of the **".npk"** installation file for the plug-in (on the installation CD or the directory to which the file was downloaded from the website).

5 Select the platform-specific binary file for the plug-in. The binary file for Windows is named **"win-x-x-x-x.npk"** and the binary file for Linux and UNIX is named **"nvf-x-x-x-x.npk"** (where **x-x-x-x** represents the version, build, and platform numbers).

   Click **Install Plugin** to begin installation.

6 After the plug-in is installed successfully, a message is displayed.

Removing the plug-in

To remove the plug-in:

1 In the Navigation pane, click **Manage Clients**.

2 In the **NetVault Clients** list, select the client, and click **Manage.**

3 In the Installed Plug-ins table, select **File System**, and click the Remove Plugin button ( ).

4 In the confirmation dialog box, click **Remove**.

Using the correct plug-in binary on Solaris

On Solaris, the **startup.sh** script sets the binary to use for the NetVault.

NetVault uses the **startup.sh** script to determine the mode the system is running in, and accordingly run the 32-bit or 64-bit binary for the NetVault.

To run the 32-bit or 64-bit plug-in based on OS mode, use the following command to start NetVault on Solaris:

```
startup.sh start
```

3

Configuring the plug-in

- Configuring default settings
- Exclusion lists

Configuring default settings

To configure the default settings for NetVault:

1 Start the Backup Job Wizard, and click **Create New** next to the **Selections** list. Open the NetVault Server or Client node. Select **File System**, and click **Configure**.

You can also configure the default settings from the **Change Settings** page. In the Navigation pane, click **Change Settings**. On the **Configuration** page, click the **Server Settings** or **Client Settings**, and then on the Settings page, click **Plugin Options**.

2 Under **File System Plugin**, configure the following settings.

Option	Description
Display Mapped Drives under Network Share	This check box is selected by default. If you do not want to display mapped network drives under the **Network Share** node, clear this check box. This option is applicable only to Windows-based clients.
Exclusion List Settings	An Exclusion List is a text file that contains a list of files and directories that you want to exclude during a backup or restore job. The NetVault supports name-based and path-based exclusion lists. Exclusion lists are added to a client, and apply to all backup and restore jobs performed on that client. For more information about exclusion lists, see Exclusion lists. To specify an exclusion list for a backup or restore job, type the file path for the exclusion list in the appropriate box: • **Backup Exclusion List — Name Based Patterns:** Name-based exclusion list for backup jobs. • **Backup Exclusion List — Path Based Patterns:** Path-based exclusion list for backup jobs. • **Restore Exclusion List — Name Based Patterns:** Name-based exclusion list for restore jobs. • **Restore Exclusion List — Path Based Patterns:** Path-based exclusion list for restore jobs. To remove an exclusion list, clear the corresponding box.
Backup Windows File Short Names	To back up short name attribute for selected files, select this check box.
Restore Windows File Short Names	To restore short name attributes for selected files, select this check box. This option is only applicable to Windows clients. To use this option, you must select the **Backup Windows File Short Names** option during backup.

Option	Description
Memory Cache Settings for Analyzing Directories and Entries	The memory cache settings affect the “restore preparation phase.” During this phase, the plug-in analyzes the saveset and examines the index to determine what it needs to restore. This information allows the plug-in to read a saveset efficiently and restore directory structures for the selected entries. The plug-in retrieves the target items from the saveset and stores them in a temporary file that is read in the memory. If the list is too large (for example, includes millions of files), the entries are moved to the disk as the items are processed. The memory cache options control the number of items that can be read from the disk at any one time. If more entries are read at any one time and more memory is allocated to store them, the speed of processing (and thus the speed of restore) increases. When the restore includes many files and directories (for example, millions of files) or deep directory structures, it is beneficial to increase the memory cache for the directories and entries. These settings improve the accessing of the index file and not the actual data transfer rate. The memory settings are specific to an individual client, and include the following: • **Initial Size of Memory Cache when Analyzing Directories to Restore:** The minimum memory cache available to the plug-in for analyzing directories that you want to restore. The default value is 1MB. • **Maximum Size of Memory Cache when Analyzing Directories to Restore:** The maximum memory cache available to the plug-in for analyzing directories that you want to restore. The default value is 10MB. • **Initial Size of Memory Cache when Analyzing Entries to Restore:** The minimum memory cache available to the plug-in for analyzing entities that you want to restore. The default value is 1MB. • **Maximum Size of Memory Cache when Analyzing Entries to Restore:** The maximum memory cache available to the plug-in for analyzing entities that you want to restore. The default value is 10MB.
Cluster Backup Directory	In a Windows cluster setup, specify the shared location to store the cluster database backup file. The format to specify the directory is: `\\<machine name>\<share name>`
Delete Cluster Backup File From Disk after Backup	After a cluster backup completes, the plug-in deletes the cluster database backup file copied to the shared location. If you do not want to delete the database backup file, clear this check box.
Reset Access Times	By default, the plug-in does the following during backups: • Updates the Access Time attribute for all files that are backed up. • Uses the Change Time attribute to select files and directories for Incremental and Differential Backups. When you select the **Reset Access Times** check box, the plug-in does the following: • Resets the Access Time attribute for all files included in a backup after the file is accessed. • Uses the Modify Time attribute to select files and directories for Incremental and Differential Backups. **NOTE:** On Linux and UNIX systems, resetting Access Time updates the Change Time attribute. Therefore, the plug-in uses Modify Time for Incremental and Differential Backups instead of Change Time.

Option	Description
	When Modify Time is used, the files and directories with attribute-only changes are not selected for Incremental and Differential Backups on Linux OS. Examples of attribute changes include: • Changing file mode or permissions • Changing file owner • Changing file group • Hard-linking to the file • Changing the file ACL • Changing the extended attributes These changes update the Change Time (inode ctime) and are included in a backup if the **Reset Access Times** check box is not selected. **NOTE:** Currently, the plug-in cannot reset Access Time on file systems that save file times in sub-second form. These file systems include NTFS, ext4, ZFS, and HFS (Mac OS X 10.6). However, the plug-in continues to use the Modify Time attribute instead of the Change Time for Incremental and Differential Backups.
Use VSS to Backup Registry if Available	This option allows the plug-in to use VSS to back up the registry during a System State backup. It is selected by default, and only supported on the Windows OS. To use this option, VSS must be installed and running on the target clients. When this check box is not selected or VSS is not installed or available on a client, the plug-in uses Win32 API for registry backups. **NOTE:** On Windows 2003, if a system contains many registry entries, the OS may take one to two minutes to return the Win32 API call. This delay may lock out the other processes and cause process failures and loss of network processes. Using VSS for registry backups resolves this issue. This problem may not be observed on Windows 7 or Windows 2008. Windows XP does not support the VSS Registry Writer. On this OS, the plug-in ignores this option and uses Win32 API for registry backups
Maximum File System Read	This option determines the maximum amount of data that can be read at a time. The value must be specified in KB. The default value is zero (0). If you do not want to set a maximum value for the file reads, use the default value. **NOTE:** Changing the default values for Maximum File System Read or Write options could result in a performance degradation. The backup and restore speeds may slow considerably. You should change these settings only when the network constraints demand it.
Maximum File System Write	This option determines the maximum amount for data that can be written at a time. The value must be specified in KB. The default value is zero (0). If you do not want to set a maximum value for file writes, use the default value.

3 Click **OK** or **Apply** to save the settings.

Exclusion lists

This section includes the following topics:

- About exclusion lists
- Exclusion list considerations

About exclusion lists

An exclusion list is a text file that contains a list of files and directories that you want to exclude during a job. You can use any text editor to create this list. After creating the file, save it on the client to which you want to apply the list. On Linux and UNIX, save the file in multi-byte text format. On Windows, save the file in Unicode text format.

NOTE: An exclusion list is applied to all backup and restore jobs performed on the client.

Exclusion list types

The NetVault supports two types of exclusion lists.

- **Name-based exclusion list**: Use this list to omit files and directories that match the given name pattern. To exclude a directory and its contents, the path must be terminated with a forward slash character (“/”) on Linux and UNIX and a backslash character (“\”) on Windows.

  Example

  ```
  # Name-based Exclusion List
  *.htm
  tmp/
  mp3_player.exe
  *.mp3
  install?.exe
  test*.*
  ```

  This exclusion list excludes the following directories and files during a backup:

  - All files with the extension “.htm”
  - Any directory named “tmp”
  - A specific file named “mp3_player.exe”
  - Any file with the extension “.mp3”
  - Any file named “install?.exe,” where “?” is any single character
  - All files that begin with the letters “test” (it excludes files such as “testdata.doc” and “test1.txt,” but does not exclude a file named “atest1.txt”)

- **Path-based exclusion list:** Use this list to omit files and directories in the specified path. To exclude a directory and its contents, the path must be terminated with a forward slash character (“/”) on Linux and UNIX and a backslash character (“\”) on Windows.

  Example

  ```
  # Path-based Exclusion List
  C:\Application\Data\Test\
  C:\Program Files\Microsoft\
  C:\WINDOWS\system32\
  D:\Work in Progress\
  D:\Files\*.txt
  D:\Database*\
  ```

  This exclusion list excludes the following directories and files during a job:

  - All contents of the directory C:\Application\Data\Test
  - All contents of the directory C:\Program Files\Microsoft

- All contents of the directory C:\WINNT\system32
- All contents of the directory D:\Work in Progress
- All files with the extension ".txt" in the directory D:\Files
- All directories within D:\ that begin with the characters Database

Exclusion list considerations

- The Plug-in *for FileSystem* supports the following special characters in an exclusion list.

Character	Description
/	Used to terminate Linux and UNIX file paths.
\	Used to terminate Windows file paths.
?	Used to match any one character.
!n	Used to insert new line.
!r	Used to insert carriage return.
!t	Used to insert tab.
!f	Used to insert form feed.
!b	Used to insert backspace,
!	Used to escape any special character that follows.
!!	Used to specify ! character in the filename or path.

- File and directory names are case-sensitive on Linux and case-insensitive on Windows.
- Spaces are not ignored.
- To insert comments, use # at the starting of each line.
- An exclusion list should not include any System State components on Windows, especially Windows File Protection and SYSVOL files on a Domain Controller.
- An exclusion list cannot contain files or paths with non-ASCII characters. Such entries can cause a backup to fail.
- Exclusion list is not supported for Off-Host backup.

4

Backing up data

- About backup methods
- About backup types
- About System State Backups
- About multistream backups
- Backup and recovery strategy
- Performing Full Backups
- Inclusion lists
- Additional notes
- Performing Incremental and Differential Backups
- Stopping a job
- Restarting a job
- Performing Network Share Backups
- Performing System State Backups

Backing up with Plug-in *for FileSystem*

NetVault Plug-in *for FileSystem* works with NetVault Server to back up sparse files and provide file-level recovery (FLR) of data backed up using continuous data protection (CDP).

About backup methods

The NetVault supports the following backup methods:

- **Standard Backup method:** This method lets you perform native file system backups. The Standard Backup method is supported on all operating systems.

  On Linux- and UNIX-based systems, only the Standard Backup method is available for backing up the file system data.

- **Snapshot-Based Backup method:** This method lets you perform snapshot-based backups. A snapshot provides point-in-time copy of data. Snapshots enable consistent backups of data and minimize the time an application is in backup mode.

  You can use the Snapshot-based Backup method to perform the following types of backups:

  - Use snapshots to perform consistent backups to disk- or tape-based storage devices.
  - Create and store snapshots as backups on supported disk arrays.

  The snapshot-based method is only supported on Windows-based clients. For more information about supported OS versions and disk arrays, see the *Quest NetVault Compatibility Guide*.

NOTE: The Backup Files to Storage is supported on any Windows platform and any disk storage. To use the **Array-Based Persistent Snapshot** and **Discard After** the data that you want to back up must reside on a supported disk array. Also, for persistent snapshot only backup, only metadata is copied to the target device and snapshot is created on storage array.

About backup types

The NetVault supports the following backup types:

- **Full Backup:** A Full Backup provides a backup of all selected files and folders. Full Backups take longer to complete and consume more backup media. However, restores are faster since the plug-in is only required to restore a single saveset. Full Backups serve as the base backups for subsequent Incremental and Differential Backups.
- **Incremental Backup:** An Incremental Backup provides a backup of files that are new or have changed since the last Full or Incremental Backup. Incremental Backups consume minimum storage space and are quicker to perform. However, data recovery can take more time depending on the number of savesets that the plug-in is required to restore.
- **Differential Backup:** A Differential Backup provides a backup of files that are new or have changed since the last Full Backup. Differential Backups speed up recovery since the plug-in is only required to restore two savesets. However, these backups consume more storage space and take more time to complete than the Incremental Backups. Differential Backups duplicate the data backed up by previous backups of the same type.

About System State Backups

System State Backups back up critical system data, which is essential to recover a functional system in the event of a system failure. These backups are only supported on Windows-based systems. System State data include the following components:

- Registry
- COM+ Class Registration Database
- Boot files, including the system files
- System files that are under Windows File Protection
- Certificate Services Database (on a Certification Services Server)
- Active Directory (on a Domain Controller)
- SYSVOL (on a Domain Controller)
- Cluster Service Information (in a cluster setup)
- IIS Metadirectory (if it is installed)

The Snapshot-Based Backup method is recommended for a consistent backup of the System State data. Regardless of the backup type, certain System State components are always backed up in full. You must include all system state components in a backup or restore job.

NOTE: In a cluster environment, you may require the Disk Signature information during a System State restore if the shared cluster disks are changed after a backup. The Disk Signature information can be obtained using the Windows **"dumpcfg.exe"** utility. For more information about this tool, consult the relevant Windows documentation.

About multistream backups

The NetVault can be configured to perform multistream operations when you are backing up large amounts of data. Multistream backups improve backup performance by using parallel streams to transfer data.

During multistream backups, the plug-in splits the Backup Selection Set is into smaller selection groups and uses separate child processes to perform parallel backups of these selection groups.

You can use the multistream options to do the following:

- Back up multiple volumes, network shares, and mount points simultaneously
- Back up a single volume using multiple parallel streams

You can specify the maximum number of parallel data streams that the plug-in can use for each job. The plug-in allows up to eight parallel streams.

NOTE: For multistream backups, when using the Use Heuristics option, granularity for selection groups is at the directory level. If you do not select Use Heuristics, granularity is at the volume level.

Backup and recovery strategy

A well-designed backup strategy helps you to recover from the damages caused by failure and resume normal operations as quickly as possible. Before you start backing up your data, ensure that a good strategy is in place that safeguards your data against various failure modes, such as media failure, data corruption, user error, and complete loss of a data center.

Typically, your backup plan should define what backup methods are used, when and at what intervals the backups are performed, how the backups are stored, how long the backups are retained, and how the backup media are re-used.

The following are some example backup sequences to offer you guidelines:

- **Full Backups only:** You can choose to perform only Full Backups if the backup size is small, backup window is not an issue, or storage media is not a constraint. In such scenarios, you can schedule Full Backups every night or every N hours depending on the frequency of updates.

  If a failure occurs, the plug-in is only required to restore a single saveset.

- **Full and Incremental Backups:** For quicker backups and minimum use of storage media, you can include Full and Incremental Backups in your strategy. For example, you can schedule Full Backups every Sunday and Incremental Backups every day or every N hours depending on the frequency of updates.

  If a failure occurs, the plug-in is required to restore data from the recent Full Backup and each Incremental Backup in the backup sequence. The restore might take longer if several Incremental savesets have to be restored. For example, if the failure occurs on Saturday, the plug-in is required to restore the Full Backup taken on Sunday and Incremental Backups taken from Monday through Friday.

- **Full and Differential Backups:** For quicker restores and reduced media usage, you can include Full and Differential Backups in your strategy. For example, you can schedule Full Backups every Sunday and Differential Backups every day or every N hours depending on the frequency of updates.

  If a failure occurs, the plug-in is required to restore data from the recent Full Backup and the last Differential Backup.

- **Windows System State Backups:** For System State backups on Windows, you can choose either of the following strategies:

  - Back up only the System State data
  - Back up the C Drive and System State data

  Backing up the System State data with C drive has the following advantages:

- It eliminates duplication of the system files in a saveset. If you run two separate jobs, the system files are included in both savesets.
- You can use this backup to perform a disaster recovery.

The snapshot-based method is recommended for a consistent backup of the System State data.

- **Policy-based backups:** At the enterprise-level, you can create a policy to back up the root directory on Linux and UNIX clients and all fixed drives and System State on Windows clients. For more information about backup policies, see the *Quest NetVault Administrator's Guide*.

Performing Full Backups

To perform Full Backups:

1 In the Navigation pane, click **Create Backup Job** to start the configuration wizard.

You can also start the wizard from the **Guided Configuration** link. In the **Navigation** pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Create Backup Jobs**.

2 In **Job Name**, type a name for the job. Assign a descriptive name that allows you to easily identify the job for monitoring its progress or restoring data.

A job name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. There is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

3 In the **Selections** list, select an existing Backup Selection Set, or complete the following steps to create a set:

a Click **Create New** to open the **NetVault Selections** page.

b Open the NetVault machine, and then open **File System**.

c Select the data that you want to back up:

- **Back up entire system:** On Linux and UNIX systems, select the root node ("/").

  On Windows, select the **Fixed Drives** node. To back up data on removable drives, select the **Removable Drives** node.
- **Back up individual items:** Open the parent node. Browse through the directory tree, and locate the items that you want to back up.

  To include a file or a directory, select the corresponding check box. When you select a parent node, the child items are automatically included in the backup.
- **Exclude child items:** To exclude a file or directory within a selected parent node, click the corresponding check box to replace the check mark with a cross.

d Click **Save**, and in the **Create New Set** dialog box, type a name for the set. Click **Save** to close the dialog box.

NOTE: A set name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. On Linux OS, the names can have a maximum of 200 characters. On Windows OS, there is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

4 In the **Plugin Options** list, select an existing Backup Options Set, or complete the following steps to create a set.

a Click **Create New** to open the **File System Plugin Backup Options** page.

b Select the backup method and backup type.

Option	Description
Backup Method	Under **Backup Method**, select the applicable option: • **Standard:** To perform native file system backups, select the option. On Linux- and UNIX-based systems, you can only use the Standard Backup method. • **Snapshot-Based:** To perform snapshot-based backups, select this option. This backup method is only available on Windows-based systems. The plug-in uses a VSS provider to create snapshots, and uses these copies to perform backups to disk- or tape-based storage devices. You can also use the snapshot-based method to create persistent snapshots on supported disk arrays. For more information about backup methods, see About backup methods.
Backup Type	Under **Backup Type**, select the **Full** option. For more information about backup types, see About backup types.

c On Windows-based machines, configure the following snapshot options.

Table 1. Snapshot options

Option	Description
Backup Files to Storage	This option is selected by default. To copy the selected data from snapshots to a storage device, select this check box. The plug-in uses a suitable VSS provider to create persistent or non persistent snapshot on the client, and copies the selected data from the snapshot to the storage device.
Array-Based Persistent Snapshot	To retain persistent snapshots on disk arrays, select this check box. Note the following: • If both **Backup Files to Storage** and **Array-Based Persistent Snapshot** check boxes are selected, the plug-in copies the selected files to a storage device, creates a snapshot on storage array, and adds the snapshot information to the backup index. • If only the **Array-Based Persistent Snapshot** check box is selected, the plug-in just writes the index entries to the storage device and creates a snapshot on storage array.

Table 1. Snapshot options

Option	Description
Discard After	The **Discard After** option can be used to delete persistent snapshots from Disk Array after a specified period. This option is only available when you use persistent snapshots to perform backups to a storage device and retain the snapshot as backup on the disk array (that is, the backup uses **Array-Based Persistent Snapshot** option). To configure snapshot expiration setting, select this check box. Type or select the period, and in the associated list, select the Day, Week, Month, or Year option. Regardless of the saveset expiration setting, a snapshot is automatically deleted as per its retirement settings on disk array. Otherwise, it expires as per retention policy selected while creating a snapshot. **NOTE:** In case of Huawei OceanStor Device Manager and NetApp ONTAP Device Manager, a snapshot cannot be discarded after a specific interval.
Perform Backup Off-Host using Client	This option allows you to off-host the backup on a remote client machine, where the backup is initiated by the off-host remote client. The off-host remote client mounts the snapshot, writes file data to the backup stream, unmounts the snapshot, and returns the backup status to the primary client machine. **NOTE:** In case of Huawei OceanStor Device Manager and NetApp ONTAP Device Manager, during off-host backup, the host where the LUN is mounted and the off-host client must be in the same host group for Huawei and the same initiator group for NetApp. This option is available only if **Backup Files to Storage** and **Array-Based Persistent Snapshot** are selected.

d Configure the following backup options.

Table 2. Backup options for Full Backup

Option	Description
Ignore truncated files	**NOTE:** This option is only supported on Windows-based machines and can be used with the FileSystem plug-in and Storage Manager. The Storage Manager allows running a truncation policy to move data blocks from the primary storage and free up disk space. The Storage Manager moves the files that meet the truncation criteria from the primary storage to archive tiers of storage, and sets the Windows offline attribute for these files. Depending on the truncation policy, the Storage Manager leaves a stub file on the disk when a file is truncated. If you select the **Ignore truncated file(s)** check box, the Plug-in *for FileSystem* omits the offline files during Full, Incremental, and Differential Backups. If you do not select this check box, the Storage Manager retrieves the truncated files from the archive tiers when the plug-in tries to read the stub files. The deduplicated files are rehydrated during this process. The archived truncated files will not appear during restore.
Check For Files Being Modified During Backup	To mark the files that are being modified during a backup, leave this check box selected. During restore, these files are not restored by default. To restore these files, you must set the option **Restore Files Modified During Backup** on the Restore Options Set. If you do not want to check for files that are being updated during backup, clear this check box. The **Check For Files Being Modified During Backup** option is disabled for the snapshot-based method.
Backup through Mount Points	This check box is only available on Windows. It is selected by default. If you do not want to back up the data on NTFS mount folders, clear this check box.
Backup through Remote Mounts	This check box is only available on Linux and UNIX. It is not selected by default. If you want to back up the data on mounted file systems, you must select this check box. **NOTE:** If you do not select this check box, the plug-in does not back up any remote mount points that are included in the Backup Selection Set. The plug-in excludes the remote mount points and does not provide any error or log messages in the Job Logs.
Check for and Skip Locked Mandatory Locking Files	To check whether mandatory locking is enabled on the system or not, select this check box. If mandatory locking is enabled, the plug-in tries to lock the file before backing it up. If the plug-in is unable to lock the file, it skips the file. (For example, if the file is locked by some other application, the plug-in cannot lock it for backup, and hence skips the file.) This check box is only available on Linux and UNIX systems.

Table 2. Backup options for Full Backup

Option	Description
Enable Restartable Backup	To add the ability to stop a job and resume it later from that point, select this check box. When you stop the job, the plug-in generates an index for all items that have been processed up to that point and writes the index to the backup media and NetVault Database. The job status is then set to **Job Stopped**. If the plug-in is writing a large backup index, the jobs status continues to be reported as "Writing to Media: Storing Backup Index" until the index is written. When you restart the job later, the plug-in runs an Incremental Backup job to back up the remaining files and folders. You can stop and resume a job from the **Job Status** page. For more information, see Stopping a job and Restarting a job.
Path to Backup Log	To generate a backup log file, type the filename. The log provides a list of files that are selected for the backup. The files that are backed up successfully are marked with an "o" and the others are marked with an "x." You can use this option with an Incremental Backup to determine which new or changed files have been backed up. If you specify an existing filename, the plug-in overwrites the file. The log file is created in CSV file format (.csv), and includes details such as file size, date modified, and file type.

e Configure the following options to perform multistream backups.

Option	Description
Utilise Multiple Streams During Backups	By default, the plug-in generates a single data stream for a backup job, and backs up the selected items in a sequential manner. You can use this option to perform multistream backups. For more information, see "About multistream backups." When you select this check box, the plug-in splits the Backup Selection Set into multiple selection groups. It creates a selection group for each volume and network share included in the set and uses separate child processes to perform parallel backups of the selection groups. You can also configure the plug-in to create separate selection groups for a single volume and use multiple streams to back up the data. For more information, see Use Heuristics for large Single Volume backups. You can specify the maximum concurrent streams on a per-job basis. The **Maximum Concurrent Streams** setting determines the number of selection groups that are backed up simultaneously.
Maximum Concurrent Streams	The value configured for the **Maximum Concurrent Streams** option determines the number of child processes that are started concurrently on the client. For example, if there are six selection groups, and this option is set to 4, four child processes are started on the client. After acquiring the backup device and media, a child process generates a backup stream to back up the items within the selection group. When a child process completes its task, another child process is then started to back up the next selection group. This process continues until all selection groups are backed up. If any child process is unable to acquire the backup device and media, it remains in a "waiting for media" state until the device and media become available. For example, if only two tape devices are available for four child processes, the processes that are able to acquire the device and media start backing up the data immediately, while the other processes remain in a waiting state until the device and media become available.

Option	Description
	We recommend that you configure this option according to the capabilities of the client on which the job is running and the number of devices that are available to the job. The inability to acquire device and media can delay the job. In addition, running multiple concurrent processes on systems with a single processor can degrade the system performance. By default, the plug-in uses two streams for backups that are configured as multistream jobs. The plug-in allows up to eight parallel streams. Note the following: • A single stream is used to back up all items in a selection group. • If any child process fails to back up the items in a group, the entire job fails. • If any child process reports any warnings, the entire job completes with warnings. • Restores of multi-stream backups are performed sequentially. **NOTE:** The Backup Options Sets (default and user-defined) created using a plug-in version prior to 11.0 do not include the multistream options. To apply multistream options to pre-existing jobs, edit the job and assign a set that includes these options.
Use Heuristics for large Single Volume backups	When backing up a large volume, you can configure the plug-in to dynamically split the data into multiple selection groups and use multiple parallel streams to back up these groups. The criteria used to determine the selection groups include the following: • Number of peer directories • Number of sub-directories • Number of files within a directory After identifying the selection groups, the plug-in creates a child process to back up each selection group. Each child process uses the same criteria to identify parts of its backup that can benefit from a separate stream. If the current stream count has not exceeded the maximum allowed streams, the plug-in assigns a new child process to back up the sub-group. Otherwise, it uses the current stream to back up the sub-group. During Incremental Backups, only the top-level selection groups use separate data streams; the child processes do not attempt to split the selection groups. **NOTE:** This feature is intended for single volume selections that contain large amounts of data (for example, volumes that are larger than 1TB in size). It is not likely to benefit small backups and should not be used in those cases.
Always use a new stream when crossing a mount point	When you select this check box, the plug-in uses separate streams to back up the mount points that exist on the volumes. You can use this option with the **Use Heuristics for large Single Volume backups** option.

f Click **Save**, and in the **Create New Set** dialog box, type a name for the set. Click **Save** to close the dialog box.

5 Select or create the Schedule Set, Target Set, and Advanced Options Set. For more information about these sets, see the *Quest NetVault Administrator's Guide*.

NOTE: When performing multi-stream backups, do not select the **Ensure backup is the first on the target media** check box. If you select this check box for multi-stream backups, each data stream targets a separate piece of media to exist as the first backup on the media item. Thus, if a backup generates five streams, the job tries to obtain five blank or new media items.

This option does not apply to disk-based storage devices.

NOTE: For multistream backup's granularity for the selection groups is at *Directory* level when **Use Heuristics for large Single Volume backups** option is selected (granularity is at *Volume* level when **Use Heuristics for large Single Volume backups** option is not selected).

6 To submit the job for scheduling, click **Save & Submit**. You can monitor the job progress from the **Job Status** page and view the logs from the **View Logs** page.

To save the job definition without scheduling it, click **Save**. You can view, edit, or run this job from the **Manage Job Definitions** page. This job is not displayed on the **Job Status** page until you submit it.

For more information about **Job Status**, **View Logs**, and **Manage Job Definitions**, see the *Quest NetVault Administrator's Guide*.

Selection tree icons: Linux and UNIX

Table 3. Backup selection tree icons: Linux and UNIX

Icon	Description
	Open directory
	Closed directory
	Open linked directory
	Closed linked directory
	Open local mount point
	Closed local mount point
	Open remote mount point
	Closed remote mount point
	File
	Linked file
	Socket
	Linked socket
	FIFO
	Linked FIFO

Table 3. Backup selection tree icons: Linux and UNIX

Icon	Description
	Block device
	Linked block device
	Character device
	Linked character device
	Broken link
	Link
	Files modified during backup

Selection tree icons: Windows

Table 4. Backup selection tree icons: Windows

Icon	Description
	Fixed drives
	Drive
	Open directory
	Closed directory
	File
	Compressed file
	Hidden file
	Executable file
	File modified during backup
	Removable drives
	CDROM
	Removable media
	Network shares

Table 4. Backup selection tree icons: Windows

Icon	Description
	System state
	Active Directory (Windows server only)
	Com+ class registration database or registry
	Disk quota database
	System files
	SYSVOL
	Remote storage manager database
	IIS meta database (Windows server only)
	Certificate server (Windows server only)

Inclusion lists

This section includes the following topics:

- About inclusion lists
- Inclusion list considerations
- Adding an inclusion list
- Removing an inclusion list

About inclusion lists

An inclusion list is a text file that contains a list of files and directories that you want to include in a job. You can use any text editor to create this list. After creating the file, save it on the client to which you want to apply the list. On Linux and UNIX, save the file in multi-byte text format. On Windows, save the file in Unicode text format.

NOTE: Inclusion lists are added on a per-job basis.

Inclusion list considerations

- The maximum size of an inclusion list depends on how much memory is available to the backup process while it is running. It also depends on the system load at the time of backup. The list can still contain tens of thousands of entries on a small system to hundreds of thousands of entries on a 64-bit system with large memory configuration.
- Provide the full file path for the file or directory that you want to include in a backup. An inclusion list does not support wildcards.
- Specify one inclusion per line.

- Inclusion lists cannot contain files or paths with non-ASCII characters. Such entries can cause a job to fail.
- Inclusion list is not supported for Off-host backup.

Adding an inclusion list

To add an inclusion list:

1. On the **NetVault Selections** page, open the NetVault Client to which you want to add the inclusion list.
2. In the list of plug-ins, select **File System**, and click **Show inclusion lists**.
3. Open the **Inclusion Lists** node.
4. Use one of the following methods to add the inclusion list:
    - **Select the file from the directory tree:**
        a. Open **Browse for inclusion files**.
        b. Open the applicable nodes to display the directory that contains the inclusion list.
        c. Select the check box corresponding to the Inclusion List.
    - **Specify the file path:**
        a. Select **Specify inclusion files** and click **Add inclusion file**.
        b. In the **Enter path to inclusion file** dialog box, type the file path for the inclusion list.
        c. Click **OK**.
5. Save the Selection Set.

Removing an inclusion list

To remove an inclusion list:

1. On the **NetVault Selections** page, open the applicable NetVault machine.
2. In the list of plug-ins, select **File System**, and click **Show inclusion lists**.
3. Open the **Inclusion Lists** node.
4. Depending on the method used to add the Inclusion List, open **Browse for inclusion files** or **Specify inclusion files**. Open the directory tree, if applicable, and clear the check box corresponding to the Inclusion List.
5. Save the Selection Set.

Additional notes

- **Backup file path length limitations:** The maximum length of a backup file path depends on the operating system and the file system the data resides on.
    - On Windows, the maximum length for a path is 260 characters. It includes a drive letter, colon, backslash, components separated by backslashes, and a terminating null character (for example, D:\<256 chars>NUL). Windows also supports extended-length path of up to of 32,767 characters. It includes components separated by backslashes, each up to 255 characters in length (for example, “\\?\D:\very long path”).
    - On Linux and UNIX platforms, the maximum length for a path is 1024 characters, including the target filename and path. The path can consist of a maximum of 255 elements. An element here

refers to individual directory named in the path as well as the target filename. Restores of data to platforms or file systems with different namespace support may result in truncation of some path elements. These limitations do not apply to all UNIX file systems, particularly ZFS.

For backups of Linux and UNIX clients, NetVault uses the **cpio** tape format that does not impose these limitations. However, OS imposed limitations do apply. For more information about the OS imposed limitations, see the relevant OS documentation.

- **Hard links on Linux- and UNIX-based clients:** Various instances of hard-linked files on the Linux and UNIX clients are considered separate entities. The NetVault backs up the complete file and the link information for each selected hard-linked file. To back up all instances of a hard-linked file, include each instance in the selection.

  For example, if a file /usr/data1/a.txt is hard-linked to the file /usr/data2/b.txt, selecting /usr/data1/a.txt backs up only this file (with the link information). To back up /usr/data2/b.txt, you must include this file in the backup.

  NOTE: To maintain synchronization between the files, you must back up and restore all instances of hard-linked files together.

- **Symbolic links on Linux- and UNIX-based clients:** When a symbolic link is explicitly selected for a backup, only the link information is backed up. The plug-in does not back up the file or directory the link points to.
- **Solaris zones partitioning:** On systems using the Solaris Zones Partitioning Technology, the non-global zones share files with the global zone using loopback file system read-only mounts, usually /usr, /lib, /sbin, and /platform. When the global zone (that is, "/") is selected for a file system backup, these directories are backed up for the global zone and for each zone configured on the system. For example, if four zones are configured on the system and root is selected for a file system backup, the plug-in backs up five copies of the global zone directories. To avoid multiple backups of the directories during the global zone backup, de-select them individually, or create an exclusion list
- **Special mount points:** The NetVault treats certain file system types as special mount points, and does not recurse into these mount points. However, the mount point itself, which is a directory, is backed up.

  Consider a Linux client on which the `proc` file system is mounted as follows:

  ```
  none on /proc type proc (rw)
  none on /var/named/chroot/proc type proc (rw)
  ```

  When you perform a backup of the entire system, the directories /proc and /var/named/chroot/proc are included in the backup. The plug-in tries to back up the files in the directory /var/named/chroot/proc if it is configured Bind on Linux. However, the backup job completes with warnings. To avoid this issue, you can add the directory /var/named/chroot/proc to the exclusion list.

  The special mount points on various operating systems are listed in the following table.

Table 5. Special mount points of various operating systems

Operating system	File system type
Linux and HP-UX	• proc • swap • sysfs
FreeBSD and Mac OS X	• procfs • kernfs • swap • devfs • fdesc • volfs
NCR	• proc • stats

Table 5. Special mount points of various operating systems

Operating system	File system type
OSF	• procfs
Solaris	• proc • stats • mntfs

In addition, the following file systems are also treated as special mount points:

- bind
- rbind

- **Special files:** On Solaris, doors are not backed up. On Linux and UNIX, sockets are not backed up.
- **Windows registry keys:** On Windows, the files listed under the following registry keys are not backed up:
  - `HKLM\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup`
  - `HKLM\SYSTEM\CurrentControlSet\Control\hivelist`

  If a directory included in a Backup Selection Set is changed to a mount point, the subsequent Incremental Backups using the same set are performed as Full Backups. This issue occurs even if an intervening Full Backup is performed using the same set. To correct this issue, you must re-create or edit the Backup Selection Set.
- **Distributed File System (DFS) paths**: If a user account configured to access the root node in the DFS tree does not have access to the shared folders, a symbolic link icon accompanies the shares instead of a folder icon. If the user does not have necessary privileges, a backup completes with warnings.

  The NetVault only backs up the DFS files and directories; it does not back up the DFS tree. If the tree structure does not exist during restore, the data is restored to the local file system.

  To back up the DFS tree, you can run a backup pre-script on the client that contains the DFS root to write the tree structure to a file that is included in the backup. Alternatively, you can back up the DFS structure external to NetVault. For example, you can use the Windows tool **dfscmd** with the “batch” option to write the DFS structure to a “.bat” file:

  ```
  dfscmd /view <\\rootservername\namespace /batch > <path\filename.bat>
  ```
- **Extended attributes on NSS volumes:** To back up and restore the novell.metadata extended attributes on NSS volumes, install the attr package, and enable ListXattrNWMetadata and CtimeIsMetadataModTime switches. The following example illustrates how to enable these switches from the NSS Console (NSSCON):

  1 Start a terminal session.

  2 Log in as root.

  3 Type the following commands:

  ```
  nss /ListXattrNWMetadata
  ```

  ```
  nss /CtimeIsMetadataModTime
  ```

  These settings temporary and not saved when you reboot the machine.

  To set these switches permanently across reboots, open the “nssstart.cfg” file (located in the “/etc/opt/novell/nss/” directory), and add the following lines:

  ```
  /ListXattrNWMetadata
  ```

  ```
  /CtimeIsMetadataModTime
  ```

  Verify that the switches are spelled correctly and do not have any spaces after the forward slash character (“/”). If the switch names are specified incorrectly in the “nssstart.cfg” file, parsing errors can prevent the NSS pool from mounting.

- **Cluster data:** While backing up cluster data using a virtual client, only the LUNs owned by the active node are backed up; LUNs owned by passive nodes are not backed up.
- **Sparse files:** A sparse file uses the file system more effectively by preventing it from using disk space when sections of the file are empty. The disk space is used only when it is needed. Sparse files are useful for backup utilities, disk images, etc.

  NetVault Plug-in for Filesystem creates backups of sparse files by converting them into non-sparse files. The size of the backup will be the total size of the sparse file and not its size on the storage disk. For example, if a sparse file has a total size of 512MB and its size on the disk is 4KB, NetVault will back up the total size of 512MB. To reduce required storage space, NetVault relies on backup device compression; for example, Quest QoreStor's advanced source-based data deduplication and compression technologies are used if you are using QoreStor as a backup device.

  NetVault Plug-in for Filesystem restores the sparse files as a sparse, except for the last block. If the block size is 4KB, the size on the disk of the restored sparse file would be about 4KB more than the size on the disk of the backed-up file. For example, if the size on the disk of the backed-up sparse file is 4KB, the size on the disk of the restored sparse file would be 8 KB.

Performing Incremental and Differential Backups

To perform Incremental or Differential Backups:

1 Start the backup job wizard.

2 In **Job Name**, specify a name for the job.

  Assign a descriptive name that allows you to easily identify the job for monitoring its progress or restoring data. The job name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. There is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

3 In the **Selections** list, select the Backup Selection Set that was used to perform the Full Backup.

4 In the **Plugin Options** list, select an existing Backup Options Set, or complete the following steps to create a new set:

  a Click **Create New** to open the **File System Plugin Backup Options** page.

  b Configure the following settings.

Option	Description
Backup Method	Under **Backup Method**, select the applicable option: • **Standard:** For native file system backups, select the option. It is the only available method on Linux and UNIX systems. • **Snapshot-Based:** On Windows systems that support VSS, select this option to perform a snapshot-based backup. On Linux and UNIX systems, this option is not available. For more information about backup methods, see About backup methods.
Backup Type	Under **Backup Type**, select the applicable option. For more information about backup types, see About backup types. • **Incremental:** Select this option to back up files that were created or modified since the last Full, Incremental, or Differential Backup. • **Differential**: Select this option to back up files that were created or modified since the last Full Backup. • **Create Dump Type Backup:** Select this check box to create Dump Type Incremental or Differential Backups. With these backups you can only restore those files that were backed up in a particular Incremental or Differential Backup. **NOTE:** Regular Incremental Backups are linked to the initial Full Backup and previous Incremental Backups. Before creating a data copy of a regular Incremental Backup, you must have created a copy of the initial Full Backup and the previous Incremental Backups in the backup sequence; otherwise, the Plug-in *for Data Copy* reports an error. Dump Type Incremental Backups are standalone Incremental Backups. Such backups can be copied without creating a copy of any other backup in the backup sequence.
Ignore Archive Bit	On Windows, the archive bit helps to determine whether a file has been modified since the last backup. If the file has changed, the archive bit is set to 1, and if the file has not changed, the bit is set to 0. By default, the plug-in backs up a file for which the archive bit is set to 1, and clears the bit after the backup completes. Regardless of the state of the archive bit, the plug-in backs up a file if certain information, such as the file size or the last modified time, has changed since the previous backup. **NOTE:** The archive bit is not a clear indication that a file has changed. It only indicates that the bit has been set. Some applications intentionally modify the archive bit in a way that is not consistent with changes to the file. To ignore the archive bit settings and use the file size, last update date, and other file attributes to back up a file, select the **Ignore Archive Bit** check box. This option can be useful in the following events: • A program running on the system is changing the state of the archive bit, causing the plug-in to back up unnecessary files and folders. • The plug-in is able to access the files for a backup, but cannot clear the archive bit settings after the backup completes because of access rights violation. When you select the **Ignore Archive Bit** check box, the plug-in omits the files and directories for which only the file attributes such as Discretionary Access Control List (DACL), System Access Control List (SACL), and Owner Attribute have changed. If you want to back up such files, do not select this check box.

c Configure the additional options that you want to use. For more information, see Backup options for Full Backup.

d Click **Save**, and in the **Create New Set** dialog box, type a name for the set. Click **Save** to close the dialog box.

5 Select or create the Schedule Set, Target Set, and Advanced Options Set. For more information about these sets, see the *Quest NetVault Administrator's Guide*.

6 To submit the job for scheduling, click **Save & Submit**. You can monitor the job progress from the **Job Status** page and view the logs from the **View Logs** page.

To save the job definition without scheduling it, click **Save**. You can view, edit, or run this job from the **Manage Job Definitions** page. This job is not displayed on the **Job Status** page until you submit it.

For more information about **Job Status**, **View Logs**, and **Manage Job Definitions**, see the *Quest NetVault Administrator's Guide*.

Stopping a job

The NetVault lets you stop a job at any point and resume it later from the same point. To have this ability, you must configure the job as restartable

For more information about this option, see the Performing Full Backups.

When you stop the job, the plug-in generates an index for all items that have been processed up to that point and writes the index to the backup media and NetVault Database. The job status is then set to **Job Stopped**. If the plug-in is writing a large backup index, the jobs status continues to be reported as "Writing to Media: Storing Backup Index" until the index is written. When you restart the job later, the plug-in runs an Incremental Backup job to back up the remaining files and folders.

NOTE: The **Stop** and **Restart** methods do not work if you select multiple job instances simultaneously.

To stop a job:

1 In the Navigation pane, click **Job Status**.

2 In the list of jobs, select the job, and click **Stop**.

3 In the confirmation dialog box, click **OK**.

NOTE: The stop action is not immediate. The job may continue for some time after you confirm that you want to stop backing up.

Multi-stream backups that are in progress may not be affected by the Stop action, as stopping a backup does not interrupt streaming data. Any data submitted to a stream during the backup job continues to back up. After the in-progress stream is complete, no new streams begin.

Restarting a job

The Restart method lets you resume a backup from the point the job was stopped. To have this ability, you must configure the job as restartable. When you restart a job, the plug-in runs an Incremental Backup job to back up the remaining files and folders.

NOTE: The **Stop** and **Restart** methods do not work if you select multiple jobs simultaneously.

To restart a job:

1 In the Navigation pane, click **Job Status**.

2 In the list of jobs, select the job, and click **Restart**.

Performing Network Share Backups

Currently, the support for network share backups is limited to Windows network shares. Several problems might be encountered if you use this feature to back up CIFS-compatible shares, such as Samba, Snap Appliance Filers, or Mac OS X Shares. The problems include the inability to back up files with certain filename encoding or lengths, changing the filename while backing them up, and not restoring the permissions correctly.

The procedure for performing network share backups includes the steps outlined in the following sections:

- Prerequisites
- Configuring a network share
- Configuring a login account for network share
- Backing up network share data

Prerequisites

Before you start a Network Share Backup, verify that the following requirements are met:

- **Use the "soft" mount option:** When backing up NFS shares, use the "soft" mount option for mount points.

  When you use the "soft" mount option, the Plug-in *for FileSystem* terminates an operation and reports an error when an NFS request is timed out. Thus, if the plug-in is trying to access a stale mount point, the operation is terminated after the request is timed out.

  With the "soft" mount option, you can set the **timeo=<value>** and **retrans=n** options:

  - The **retrans** option determines the number of times a request is retried before an error is reported. The default value for **retrans** is three.
  - The **timeo** option specifies the number of seconds to wait before a retransmission is sent. The default value for **timeo** may vary depending on the OS, but it is typically 600 seconds.

  After setting these options, we recommend that you perform a test by browsing a stale mount point (use the command `ls <mount point>`) to verify that the request is timed out after the specified interval.

  For more information about the mount command and the available options, see the relevant OS documentation.

  When you use the "hard" mount option, an NFS request is tried indefinitely, causing the plug-in to wait indefinitely.

  If you use the "hard" mount option and the machine exporting the NFS share is not on the network, the plug-in job accessing the "stale" mount point stops responding. When this issue occurs, kill the nvfs process using `kill -9` or end the process from the Windows Task Manager. The Job Manager Process associated with the job terminates after the timeout period (usually 300 seconds). Until this timeout has elapsed, you cannot use the **Run now** method to run the job because the instance is considered active until the Job Manager exits.
- **Grant "change" privileges for Incremental Backups:** When performing Incremental Backups of network shares, the user that has been configured to log on to a particular share must be given "change" permissions to the shared folder. If you do not grant this permission, the job instance repeatedly backs up the modified files because the plug-in is unable to update the file attribute to indicate that the file has been backed up.

Configuring a network share

To back up a shared folder on a non-NetVault Client, you must configure a network share on the NetVault Client from which you want to perform the backup job.

To configure a network share:

1. Start the Backup Job Wizard, and click **Create New** next to the **Selections** list.
2. On the **NetVault Selections** page, open the applicable client, and then open **File System**.
   - Select **Network Shares** and click **Add Network Share**.
3. In the **Add New Network Share** dialog box, specify file path for the shared drive in the following format:

   `\\<IP Address or Resolvable Network Name>\<Share Name>\`

   NOTE: It is not recommended to execute Backup and Restore jobs on the hidden share locations. For example, `\\<IP Address or Resolvable Network Name>\C$\`, where `C$` is the hidden share location.
4. Click **Add** to add the network share and close the dialog box.

Configuring a login account for network share

To configure a user account to log on to the network share:

1. On the **NetVault Selections** page, select **Network Shares**, and click **Connect As**.
2. In the **Connection Details** dialog box, provide the following information:
   - **Domain:** Type the Windows Domain name for the system that contains the shared folders.
   - **Username:** Type a Domain Administrator user name. The network share backups must be performed using a Domain Administrator account to fully retrieve the file and directory permissions during a restore. A user belonging to the Administrators Group does not have the Domain Administrator privileges.

     NOTE: If a Non-Domain Administrator account is configured for backups, the file and directory permissions must be set manually after a restore.
   - **Password:** Type the password for the user account.
3. Click **OK** to save the connection details and close the dialog box.

Backing up network share data

To back up network share data:

1. On the **NetVault Selections** page, double-click **Network Shares**.
2. Select the data that you want to back up.
   - **Back up all network shares:** Select the **Network Shares** node.
   - **Back up individual shared folders:** Select the check boxes corresponding to the items that you want to back up.
   - **Back up individual items:** Open the parent node. Browse through the directory tree, and locate the items that you want to back up.

     To include a file or a directory, select the corresponding check box. When you select a parent node, all the child items within the node are automatically included in the backup.

- **Exclude child items:** To exclude a file or directory within a selected parent node, click the corresponding check box to replace the check mark with a cross.

Click **Save** to create a Backup Selection Set.

3 On the **Backup Job Wizard** page, specify a name for the job.

4 Complete Step 4 through Step 6 in Performing Full Backups.

5 To perform Incremental or Differential Backups, see Performing Incremental and Differential Backups

Performing System State Backups

To back up System State data:

1 Start the Backup Job Wizard.

2 In **Job Name**, specify a name for the job.

3 Click **Create New** next to the **Selections** list, and then on the **NetVault Selections** page, open **System State**.

4 Select the data that you want to back up:

- **Back up System State data:** To back up only the System State data, select System State.
- **Back up for disaster recovery**: To use the backup for disaster recovery purposes, select **Fixed Drive** or **C Drive** and **System State**.

NOTE: The Plug-in *for FileSystem* lets you select individual System State components, but we recommend that you back up all System State components together.

IMPORTANT: On Domain Controllers, you must always include **C Drive** in System State backups and restores. When you include the **C Drive**, all information about the Group Policy Objects (GPOs) is included during backups and restores.

If you back up or restore only the System State data, you cannot edit the GPOs that are restored from the backups. When you try to edit the restored GPO, the following error message is displayed:

"Failed to open the Group Policy Object. You may not have appropriate rights."

5 Complete Step 4 through Step 6 in Performing Full Backups.

System State backups: Additional notes

- Only the SYSVOL and System Files components can be truly included in an Incremental or a Differential Backup. The remaining components are always backed up in full regardless of the backup type you select. Similarly, when you perform a backup consolidation using the Plug-in *for Consolidation*, the new reference point only applies to the SYSVOL and System Files components.
- On Windows Server 2008 and Windows 2008 R2, a fault in the VSS can cause the plug-in to randomly omit files from a System State Backup. This issue occurs when the ImagePath property for a service entry in the Windows Registry uses a forward slash character ("/") or a relative path in the file path. This issue may also be observed in Windows 7 and Windows Vista.

  To resolve this issue, you must identify and correct the invalid entries in the Windows Registry. You can use the Microsoft System Information (Msinfo32.exe) Tool and Windows Registry Editor (regedit.exe) for this purpose.

To correct the registry entries:

1. Click **Start** on the taskbar. Point to **Programs**, point to **Accessories**, point to **System Tools**, and then click **System Information**.

   Alternatively, click **Start** on the taskbar, and select **Run**. In the **Run** window, type **msinfo32**, and then click **OK**.

2. Under **System Summary**, click **Software Environment**, and then select **Services**.
3. Examine **Path** on the right-hand pane.
4. After identifying the invalid entries, open the Windows Registry Editor (regedit.exe) to edit the ImagePath.

   To open the registry editor, click Start on the taskbar, and select Run. In the Run window, type regedit, and then click OK.

5. In the **Registry Editor** window, locate the invalid entries. These entries are usually located under the key **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services**.

NOTE: The ImagePath property stores the path to the executable program for a service. The path to the executable program can also be viewed on the **General** tab of the **Services** console which can be accessed from **Control Panel > Administrative Tools > Services**.

6. Replace the forward slash character (“/”) with the backslash character (“\”), and change the relative paths to absolute paths.

- On Windows 2008, if a System State Backup fails with the error “Failed to add extra items to the selection tree,” it could be due to the following reasons:
  - The VSS system writer is missing, modified, or in an unstable state.
  - The plug-in does not have sufficient permissions for enumerating or accessing the VSS writer.
  - The plug-in is forced to use the VSS writer provided by an application, such as SQL Server or Exchange Server.

  For more information about this issue and possible solutions, see the KB article https://support.quest.com/kb/SOL81588.

5

Restoring data

- Restoring data using the plug-in
- Searching for files in savesets
- Restoring an item using the Quick Save feature
- Viewing media list
- Renaming and relocating data during restore
- Restoring data to an alternate client
- Restoring System State Backups

Restoring data using the plug-in

To restore data using the plug-in:

1 In the Navigation pane, click **Create Restore Job**.

On the **Create Restore Job — Choose Saveset** page, the saveset table provides a list of available savesets. The table shows the saveset name (Job Title and Saveset ID), creation date and time, saveset size, and saveset status.

The saveset status is indicated using the following icons.

Icon	Description
	Saveset is online (all segments are online).
	Saveset is partially online (some segments are online).
	Saveset is offline (all segments are offline).

The saveset list is sorted alphabetically by saveset name. You can sort the list by a different column or reverse the sort order by clicking the column heading. The arrowhead next to the column name indicates the sort order.

You can click **Load more** to load the next set of records, if available. Each load action fetches a maximum of 5000 records. This button is disabled if there are no more records to load.

By default, the table is sorted by creation date and time. You can sort the table by one or more columns, as required.

NetVault web application supports multiple-column sorting for the tables. However, the first sorted column takes the highest precedence over the sorting of other columns. For example, if the table is sorted by the ‘Start Time’ as the primary sort, selecting sort options on the ‘ID/ Instance/ Phase’ column does not affect the list order.

To sort by a column, click the column header; to change the sort direction, click the column header again. Repeat this step for other columns by which you want to sort the table.

The arrowhead next to the column name indicates the sorting order (up for ascending order and down for descending order). A blue arrowhead is used to indicate the sort order for the primary column.

To remove a sort, click the column header for the corresponding column until the arrowhead is no longer displayed.

To choose the following actions in a column, click the column menu:

Option	Description
Sort Ascending	Sorts the column in ascending order.
Sort Descending	Sorts the column in descending order.
Columns	Select/ clear to show/ hide the required column (s).

To set the filter options, view the page size setting, sorting order in the table, export the records, or edit table settings, click the icons at the lower- right corner of the table.

For more information, see the topic on customizing tables in the NetVault WebUI in the *Quest NetVault Administrator's Guide*.

2 Select the saveset that you want to use, and click **Next**.

When you select a saveset, the following details are displayed in the **Saveset Information** area: Job ID, job title, tag, server name, client name, plug-in name, saveset date and time, retirement setting, Incremental Backup or not, archive or not, and saveset size.

3 On the **Create Selection Set** page, select the data that you want to restore:

- **Restore entire saveset:** Select the primary or the root node.
- **Restore individual items:** Open the parent node. Browse through the directory tree, and locate the files that you want to restore. To include a file or a directory, select the corresponding check box. When you select a parent node, all the child items within the node are automatically included in the restore job.
- **Exclude child items**: To exclude a file or directory within a selected parent node, click the corresponding check box to replace the check mark with a cross.

NOTE: To restore data from a persistent snapshot, no additional steps are required. The plug-in automatically obtains the snapshot information from the backup index. If the snapshot is available, the plug-in mounts it locally and restores the selected data from the snapshot. If the snapshot is unavailable, the plug-in restores the data from the storage media. For persistent snapshot only, restore fails if snapshot is unavailable.

4 On the **Create Selection Set** page, click **Edit Plugin Options**, and configure the following options.

Table 6. NetVault Restore Options

Option	Description
Overwrite Newer Files	By default, the plug-in overwrites existing files on the restore destination with the backed-up version in the saveset. If you do not want to overwrite the existing files, clear this check box.
Reset File Timestamps	By default, timestamps for restored files are reset to what is recorded in the backup saveset. To restore the files with current timestamp, clear the check box.
Reset Directory Timestamps	By default, directories are restored with current timestamp. To reset the timestamp to what is recorded in the saveset, select this check box. To use this option, you must also select the **Reset File Timestamps** check box.
Attempt to Restore Hard Links	To re-establish the links when you restore hard-linked files, select this option. This check box is only available on the Linux and UNIX OS.
Restore Files Modified During Backup	To restore the files that were marked "in flux" during backup, select this check box.

Table 6. NetVault Restore Options

Option	Description
Path to Restore Log	To generate a restore log file, type the filename. The log file provides a list of files that are selected for a restore. The files that are restored successfully are marked with an “o” and the others are marked with an “x.” If you specify an existing filename, the plug-in overwrites the file. The log file is created in CSV file format (.csv), and includes details such as file size, date modified, and file type.
Exclusion List Path	To use an exclusion list, type the full path to the file. For more information about exclusion lists, see Exclusion lists.
Restore Meta Data	To restore the metadata on Linux and UNIX systems, select the appropriate option: • **Compatible:** By default, the plug-in restores the metadata only if it is compatible with the current file system. On Mac OS X, at least one of the following requirements must be met to use this option: Both the backup source and the restore target should be running Mac OS X 10.4.0 or later. Both the backup source and the restore target should be running the same major version of Mac OS X. For example, you can restore backups from Mac OS X 10.3.2 to 10.3.5. • **Force**: Select this option to force the plug-in to restore the metadata regardless of its compatibility with the current file system. On Mac OS X, the following requirements must be met to use this option: - The backup should be from a Mac system. - The backup should not be old. • **Discard:** If you do not want to restore the metadata, select this option.

5 Click **OK** to save the settings, and then click **Next**.

6 On the **Create Restore Job** page, specify a name for the job. Assign a descriptive name that allows you to easily identify the job for monitoring its progress.

A job name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. There is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

7 In the **Target Client** list, the client from which data was backed up is selected by default. Do not change this setting.

8 Select or create the Schedule Set, Restore Source Set, and Advanced Options Set. For more information about these sets, see the *Quest NetVault Administrator’s Guide*.

9 Click **Submit** to submit the job for scheduling.

You can monitor the job progress from the **Job Status** page and view the logs from the **View Logs** page. For more information about these functions, see the *Quest NetVault Administrator’s Guide*.

Searching for files in savesets

The **Search** option on the **Create Restore Job — Choose Saveset** page allows you to find specific files or data items without opening any savesets or browsing through their contents. You can use filenames or regular expressions to find the data items that you want to restore.

To configure or enable the catalog search, click the bulb icon shown near the **Search** button on the **Create Restore Job — Choose Saveset** page. The catalog search supports the regular expression syntax used by Elasticsearch. For more information on Elasticsearch, see

https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl-regexp-query.html. For more information on catalog search, see the *Quest NetVault Administrator's Guide.*

To search for data items in savesets:

1. On the **Create Restore Job — Choose Saveset** page, click **Search**.
2. In the **Search for files in savesets** dialog box, configure the following options:
   - **Search String:** Type the search string.
   - **Regular expression search:** To use POSIX (Portable Operating System Interface for Unix) regular expressions instead of Elasticsearch in the **Search String** box, select this check box.
   - **Use legacy search method**: This check box is displayed if both cataloged and non-cataloged savesets are included in the search.

     If only non-cataloged savesets are included in the search (or if "Use legacy search method" option is selected), the legacy search is used.

     If only cataloged savesets are included in the search (or if "Use legacy search method" option is not selected), the catalog search is used.
3. To search the data items in one or more specific savesets, select the desired saveset (s) and click **Search**. If you do not select any saveset, then all savesets are included in the search. On the **Search Results** page, you can view the specified files or data items. The following search results page is displayed when catalog search is configured and enabled on your NetVault system.
4. Select the items you want to restore. You can only restore items from one saveset.
5. Click **Restore selected items**.
6. Complete Step 3 through Step 9 in Restoring data using the plug-in.

For more information on searching for files in savesets, see the *Quest NetVault Administrator's Guide*.

Restoring an item using the Quick Save feature

Beginning with release 12.3, Plug-in *for FileSystem* users can use the Quick Save feature in the catalog search to save a particular item. There are three prerequisites for using this feature:

- Install NetVault 12.3 or later on the server.
- Configure Catalog Search on the server.
- Catalogue the backups.

To restore an item using the Quick Save feature:

1. Complete a search for the item that you want to restore by following the steps in Searching for files in savesets.

   NOTE: Quick Save applies only to files smaller than 10 MB. If you change the default file size in the search to greater than 10 MB, the default automatically resets.
2. On the results page of the catalog search, click the ellipsis (**...**) next to the item, and then select **Fetch item**.

   NetVault fetches the item from the backup image in the temporary directory under NetVault home. The request appears on the View Events page.

   An alert appears when the request is complete and the restore begins.

Viewing media list

The **Media List** option on the **Create Restore Job — Choose Saveset** page lets you view information about the media items used to store a backup. You can view details about the data segments and index segments for a backup.

To view the media details for a saveset:

1 On the **Create Restore Job — Choose Saveset** page, select the applicable saveset.

2 In the **Saveset Information** area, click **Media List**.

3 In the dialog box that appears, you can view the following details:

- **Backup size:** This area shows the total size of the saveset in number of bytes
- **Data segment table:** This table shows information about the media items that contain the data segments. You can view the following details: media label, media group label, stream ID, starting byte number, ending byte number, and media location
- **Index segment table:** This table shows information about the media items that contain the index segments. You can view the media label and media location.

4 Click **Close** to close the dialog box.

Renaming and relocating data during restore

To restore a file or directory without overwriting the existing copy, you can rename or relocate the item during restore.

To rename or relocate a file or directory:

1 Complete Step 1 through Step 3 in Restoring data using the plug-in.

2 Select the file or directory that you want to rename or relocate and click **Rename**.

You cannot rename or relocate the root directory (for example, C:\ or /).

3 In the **Rename/Relocate** dialog box, provide the following information:

- **Rename:** Type a new name for the selected item. When specifying the filename, include the extension. Otherwise, the file is restored with no extension. For example, when renaming the file “old.txt,” if you type “new” in the **Rename** box, the file is restored without the “.txt” extension.
- **Relocate**: Type the full path to the new restore location.

> **NOTE:** To rename an item without changing its location, use only the **Rename** option. To change the restore location without changing the name, use only the **Relocate** option. Use both options if you want to rename the item and change its location.

Click **OK** to close the dialog box.

After you close the dialog box, the plug-in updates the corresponding node in the selection tree to display the new name and location for the file or directory.

4 Complete Step 4 through Step 9 in Restoring data using the plug-in.

Restoring data to an alternate client

During restore, you can relocate a backup to an alternate client. This procedure can be useful during a server migration or disaster recovery operation.

Prerequisites

Before you start the restore procedure, verify that the following requirements are met:

- Install the NetVault Server or Client software on the target client. The software version must be the same as or higher than the original server or client.
- Add the target client to the NetVault Server.

Restore procedure

To relocate data to an alternate client:

1. Complete Step 1 through Step 6 in Restoring data using the plug-in.
2. In the **Target Client** list, select the restore target.

   Alternatively, click **Choose**, and in the **Choose the Target Client** dialog box, select the restore target. Click **OK** to close the dialog box.
3. Complete Step 8 and Step 9 in Restoring data using the plug-in.

NOTE: You can also use the **Rename** and **Relocate** options while restoring data to an alternate client. For more information about these options, see Renaming and relocating data during restore.

Restoring System State Backups

This section includes the following topics:

- About System State restore
- System State restore considerations
- Supported restore methods for Active Directory
- Supported restore methods for SYSVOL Directory
- Prerequisites
- Restoring System State data
- Restoring cluster database on Windows 2008
- Additional considerations for domain controllers

About System State restore

System State restore allows you to recover a functional system without reinstalling the Windows operating system. It can be useful when you want to perform the following:

- Perform a rollback (for example, overwrite the corrupted system data)
- Perform a disaster recovery

Typically, when you perform a System State restore, all components are restored to their original location. However, the NetVault also allows you to restore the System State data to an alternate directory without affecting the active system. Restoring System State in a situation other than system recovery is not recommended.

While restoring the System State data on a Domain Controller, you need to select a method for restoring distributed services like Active Directory and SYSVOL. The restore method determines whether the restored data gets replicated to other servers or not. The method you choose depends on the number and configuration of Domain Controllers in your environment.

System State restore considerations

Consider the following when restoring system state backups to a second system:

- The System State Backups can only be restored to the same system or an alternate system with the same hardware (that is, same make, model, and configuration). There should not be any driver, hardware, or firmware differences between the source and destination systems.
- On Windows 2012, Windows 2008, and Windows 2008 R2, System State recovery is only supported on the same operating system instance. It is not supported on a reinstalled operating system on the same or different hardware.
- Bare Metal Recovery (BMR) and System State Recovery have different objectives.
  - The NetVault Bare Metal Recovery Plug-ins use VSS to take a snapshot of a bootable volume (for example, C:\), and back up that volume at block level. There is no file system awareness during this process. During restore, the system is booted in an offline state using WinPE (preferred) or the Linux-based VaultOS, and the bootable volume (C:\) data is restored at block level. This method is recommended to recover a system that is in unbootable state, or to perform a full system recovery to an alternate system
  - System State Recovery recovers the System Registry, COM + Database, Certificate Services, Active Directory, SysVol, and other system state components to a previous known state. Problems can occur when you use a backup of C drive and System State to perform a full restore to an alternate system.

    For more information about these issues, see the KB article https://support.microsoft.com/en-us/help/249694/how-to-restore-a-windows-installation-or-move-it-to-different-hardware from Microsoft.

Supported restore methods for Active Directory

Microsoft supports the following methods for restoring the Active Directory on a Domain Controller:

- **Non-Authoritative Restore:** During non-authoritative restore, the distributed services on a Domain Controller are restored from the backup media, and the restored data is then updated through normal replication. Non-authoritative restore is typically performed when a Domain Controller has completely failed due to hardware or software problems.
- **Authoritative Restore:** During authoritative restore, an entire directory, a subtree, or individual objects can be designated to take precedence over any other instances of those objects on the Domain Controllers. Through normal replication, the restored Domain Controller becomes authoritative in relation to its replication partners. Authoritative restore is typically used to restore a system to a previously known state, for example, if one or more Active Directory objects were erroneously deleted.
- **Primary Restore:** Primary restore is used when the server you are trying to restore is the only running server of a replicated data set (for example, the SYSVOL and FRS).

NOTE: The NetVault supports only non-authoritative restore of the Active Directory.

Performing an Authoritative Restore of Active Directory Objects

The NetVault supports only the Non-Authoritative restore method. When you restore the Active Directory using the plug-in, the objects are restored with their original update sequence number. The Active Directory replication system uses this number to detect and propagate Active Directory changes among the other servers. The data that is restored non-authoritatively appears as old data and does not get replicated to the other servers. The Active Directory replication system updates the restored data with the newer data available on the other servers.

To perform an authoritative restore of the Active Directory data, you must run the **"ntdsutil"** utility after you have restored the System State data, but before you restart the server. The ntdsutil utility lets you mark Active Directory objects for authoritative restore. When an object is marked for authoritative restore, its update sequence number is changed so that it is higher than any other update sequence number in the Active Directory replication system. This change ensures that any replicated or distributed data that you restore is properly replicated or distributed to all servers. For more information about ntdsutil, see the relevant Microsoft documentation.

Supported restore methods for SYSVOL Directory

SYSVOL (System Volume) is a collection of folders and reparse points in the file systems that exist on each Domain Controller in a Domain. SYSVOL provides a standard location to store important elements of Group Policy objects (GPOs) and scripts so that the File Replication Service (FRS) can distribute them to other Domain Controllers within that Domain. FRS monitors SYSVOL, and if a change occurs to any file stored on SYSVOL, FRS automatically replicates the changed file to the SYSVOL folders on the other Domain Controllers in the Domain.

The NetVault supports the following restore methods for the SYSVOL directory:

- **Primary Restore:** Use this restore type only when restoring SYSVOL on a standalone Domain Controller, or on the first of several Domain Controllers. Typically, a primary restore is only required when all the Domain Controllers in the Domain are lost, and you are trying to rebuild the Domain from backup. Select Primary only for the first server. Do not use this restore type if you have already restored SYSVOL on one or more servers.
- **Authoritative Restore:** Use this restore type when you have more than one Domain Controller to roll back the SYSVOL changes, and replicate the restored data to all other servers.
- **Non-Authoritative Restore:** Use this restore type when you want to restore the data on a single Domain Controller in a replicated environment without replicating the restored data to the other servers.

Prerequisites

Before you start the restore procedure, ensure that the following requirements are met:

- Verify that there are no Automatic Updates that are applied on reboot. If necessary, disable Automatic Updates, and reboot the machine. After restoring the system state, reboot the machine, and then enable Automatic Updates. If the Automatic Updates that are installed during reboot are inconsistent with restored data, it may cause system failure.
- For a disaster recovery operation, verify that the following requirements are met:
  - The operating system version installed on the destination server is the same as the original server. All the patches and hotfixes installed on the original server are also installed on the target server.

    NOTE: All Windows updates should be performed after the restore completes.

- The number of drives on the destination server is equal to the number for drives on the original server. Each drive is of the same size as before or larger. The drive letter mapping is the same for all the drives.
- The file system format on the destination server is the same as the original server.
- The NetVault version running on the destination server is the same as the original server.

- If the destination server is a Windows Domain Controller, start it in the **Directory Services Restore Mode**. To start the server in this mode, press F8 during boot, and select **Directory Services Restore Mode** in the list of boot methods.

Restoring System State data

To restore System State data:

1. In the Navigation pane, click **Create Restore Job**.
2. In the saveset table, select the saveset that you want to use, and click **Next**.
3. On the **Create Selection Set** page, select the data that you want to restore:
   - If you only want to restore the System State data from a saveset, select the System State node.
   - If you want to perform a disaster recovery, select the C drive or Fixed Drives node, and the System State node.

   IMPORTANT: On Domain Controllers, you must always include **C Drive** in System State backups and restores. When you include the **C Drive**, all information about the Group Policy Objects (GPOs) is included during backups and restores.

   If you back up or restore only the System State data, you cannot edit the GPOs that are restored from the backups. When you try to edit the restored GPO, the following error message is displayed:

   "Failed to open the Group Policy Object. You may not have appropriate rights."
4. On the **Create Selection Set** page, click **Edit Plugin Options**.
5. Configure the standard restore options. For more information, see NetVault Restore Options.
6. Under **System State**, ensure that the ***Live Restore*** option is selected.

   NOTE: To create a copy of the System State data without affecting the active system, select the **Restore to a File** option, and type the full file path to the alternate location where you want to restore the data. This option is intended for advanced users. Quest offers no support or instructions on how to use the data restored in this manner.
7. On a Domain Controller, select the method for restoring the SYSVOL directory:
   - **Primary SYSVOL Restore:** Select this check box only when you are rebuilding the Domain. You can use this option to restore a standalone Domain Controller, or to restore the first of several Domain Controllers.
   - **Authoritative SYSVOL Restore:** Select this check box to restore the system to a previously known state, and replicate the restored data to all the other servers.
   - **Non-Authoritative SYSVOL Restore:** Select this check box to restore a single Domain Controller and update the restored data through normal replication process.
8. Click **Save** to save the settings, and then click **Next**.
9. Complete Step 3 through Step 9 in Restoring data using the plug-in.

Restoring cluster database on Windows 2008

On Windows 2008, the NetVault uses VSS for backing up and restoring the cluster database. The Windows Cluster Service must be running to use this feature. The NetVault only supports the authoritative restore schema. In an Active-Active cluster, VSS stops other nodes when you perform a restore operation. These nodes must be manually restarted after restore. On Windows 2008, the following restore options are not available:

- **Force Cluster Database Restore**
- **Use Original Cluster Database Location**
- **Cluster Quorum Drive Letter**

Additional considerations for domain controllers

- A Domain Controller in a non-authoritative state does not function as a Domain Controller until it has synchronized with the active Domain Controller. The synchronization is implemented using the Microsoft File Replication Service. If the SYSVOL data on a client has been restored non-authoritatively, the File Replication Service must successfully replicate the data to the active controller with Authoritative version of the SYSVOL data before the machine functions as a Domain Controller.
- If a restored Domain Controller owns the Flexible Single Master Operations (FSMO) roles, it does not start performing these roles until it has synchronized with the active Domain Controller. It may take some time before the restored Domain Controller starts performing the FMSO roles.

A

Troubleshooting

- Common errors

Common errors

This section describes some common errors and their solutions. It includes the following topics:

- NetVault Service fails to start on Windows
- NetVault Service fails to start after the machine is restarted
- NetVault Service starts, but stops immediately on Linux
- Snapshot-based backup fails

NetVault Service fails to start on Windows

Description

The NetVault Service fails to start on a Windows-based NetVault Server.

Symptom

Check the Windows Event Viewer to see if it displays the following message:

PDT FATAL: lock file “postmaster.pid” already exists

Solution

NetVault cannot start if the PostgreSQL database that is used to store the system data does not start. To correct this issue, delete the **“postmaster.pid”** file from the location referenced in the log and restart the NetVault Server.

NetVault Service fails to start after the machine is restarted

Description

After restarting the machine, the NetVault Service sometimes fails to start on a Windows-based NetVault Server.

Symptom

Check the Windows Event Viewer to see if it displays the following message:

FATAL: could not create any TCP/IP sockets for a PostgreSQL source

Solution

NetVault cannot start if the PostgreSQL database that is used to store the system data does not start. To correct this issue, start the Task Manager, and click **Show processes from all users**. You can see multiple instances of **postgres32.exe** running on the system. Select any one instance of this process, and click **End Process** to remove all instances of **postgres32.exe**. Then, start the NetVault Service.

NetVault Service starts, but stops immediately on Linux

Description

On a Linux-based machine, the NetVault Service starts, and then stops immediately.

Symptom

No error messages are displayed.

Solution

This issue can occur if the Postgres service cannot resolve the host name **localhost**, and fails to start.

Check the /etc/hosts file, and if the file does not contain an entry for **localhost**, add the entry.

Snapshot-based backup fails

Description

When performing a snapshot-based backup, if the VSS writer fails to generate a snapshot, the job fails.

Symptom

The log messages show the following errors:

- Failed to add extra items to the selection tree.
- Failed to prepare the selection tree and backup options.

Solution

These messages indicate a problem in the VSS subsystem. NetVault cannot complete the backup job if the VSS subsystem is not functioning properly. Before you run the job again, we recommend that you do the following:

- Delete the existing shadow copies. (You can use the **Vssadmin** or **DiskShadow** utilities to delete the shadow copies.)
- Restart the applicable VSS writer service.

About us

Quest provides software solutions for the rapidly-changing world of enterprise IT. We help simplify the challenges caused by data explosion, cloud expansion, hybrid datacenters, security threats, and regulatory requirements. We are a global provider to 130,000 companies across 100 countries, including 95% of the Fortune 500 and 90% of the Global 1000. Since 1987, we have built a portfolio of solutions that now includes database management, data protection, identity and access management, Microsoft platform management, and unified endpoint management. With Quest, organizations spend less time on IT administration and more time on business innovation. For more information, visit www.quest.com.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at https://support.quest.com.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request.
- View Knowledge Base articles.
- Sign up for product notifications.
- Download software and technical documentation.
- View how-to-videos.
- Engage in community discussions.
- Chat with support engineers online.
- View services to assist you with your product.