Quest® NetVault® Backup Plug-in *for Hyper-V* 12.0

User's Guide

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.

Contents

Introducing NetVault Backup Plug-in *for Hyper-V*

- NetVault Backup Plug-in for Hyper-V: at a glance

- Key benefits

- Feature summary

- Target audience

- Recommended additional reading

NetVault Backup Plug-in *for Hyper-V*: at a glance

Quest® NetVault® Backup Plug-in *for Hyper-V* (Plug-in *for Hyper-V*) uses a web-based user interface (WebUI) and automated workflow process to support a centralized way for configuring backup and restore policies for virtual environments. The plug-in offers a detailed level of control that minimizes downtime by letting you restore complete virtual machine (VM) images or individual files. Through integration with a range of backup devices, your data is protected and stored offsite to meet your disaster-recovery and business-continuity goals.

Key benefits

- **Increase confidence when deploying the plug-in**: Plug-in *for Hyper-V* is flexible enough to account for many recovery scenarios. By relying on the plug-in to implement your backup policies, you can focus on more critical tasks without risking your ability to recover what is needed if a failure occurs.
  Plug-in *for Hyper-V* also empowers you to create a comprehensive and flexible backup policy and reduces the need for scripting, which may increase the risk of syntax or human errors.

- **Reduce risk with flexible backup and recovery options**: Plug-in *for Hyper-V* equips administrators with a powerful set of tools to simplify backup and recovery of business-critical VMs. The Volume Shadow Copy Service (VSS)-based backup provides a simple method to protect VMs including their applications. Full image-level backups protect the VM from failures while also providing flexibility to support test environments and disaster recoveries. Individual-file recovery offers faster restores and reduces storage utilization, server load, and the amount of data transferred over the network. Plug-in *for Hyper-V* offers the choice to restore a complete VM image or individual files.

- **Shorten backup windows and improve device usage:** Plug-in *for Hyper-V* speeds VM image-level backups with high-performance Incremental and Differential Backups for data stored in VMs of Windows Server 2016. For Windows Server 2016 VMs, the plug-in uses Resilient Change Tracking (RCT). RCT tracks changes to a Virtual Hard Disk (VHD or VHDX) that occur between backups.

- **Maximize information technology (IT) efficiency**: Plug-in *for Hyper-V* reduces administrator involvement in the recovery process while delivering faster and more reliable recoveries. Less-experienced personnel can initiate restores, thus reducing downtime and improving business continuity and freeing administrators

from being available 24x7. Less-experienced personnel can also perform routine administrative options, thus empowering administrators to create comprehensive backup policies with a reduced need to understand Plug-in *for Hyper-V* internals. The plug-in uses an automated workflow with point-and-click options for tasks such as defining backups and submitting job scheduling.

- **Ensure business continuity**: With offsite backups being an important part of the data-protection plan for business-critical applications, the plug-in takes advantage of NetVault Backup's integration with a range of backup devices. Plug-in *for Hyper-V* gives you the confidence that your virtualized environments are protected and stored offsite for disaster-recovery purposes.

Feature summary

- Hyper-V Clustering
- VSS-based VM backups
- Single plug-in for image-level backups and restores
- File-level restores from image-level backups. Includes support for the following types of partitions and file systems:
  - Master Boot Record (MBR)
  - GUID Partition Table (GPT)
  - Windows Server New Technology File System (NTFS)
  - Extended file systems: EXT2, EXT3, and EXT4
  - Extents File System (XFS) on Linux and UNIX
  - Volumes managed by Logical Volume Manager (LVM) on Linux-based systems and Logical Disk Manager (LDM) on Windows-based systems as single or spanned disks
- Full support for VMs that reside on a Server Message Block (SMB) file server
- Support for NetVault Backup Virtual Client
- Support for inclusion and exclusion configuration of VM names for use with backup selection
- Support for policy-based backups
- Support for backups and restores of VMs that reside on Windows mount points in a standalone or clustered deployment.
- Uses Active Block Mapping (ABM) to reduce network and storage requirements. Includes support for the following types of partitions and file systems:
  - MBR
  - GPT
  - NTFS
  - EXT2, EXT3, and EXT4

  > **i** | **NOTE:** If you disable ABM, the plug-in automatically uses zero-block exclusion to ensure that the zeroed blocks are excluded during backups.
  > ABM is not supported with XFS.

- Restores of complete VM image or individual files
- RCT-based image-level backups of VMs that use the Hyper-V role in Windows Server 2016.
- Optional automatic power-on setting for VMs after recovery of image-level backups.
- Rename a VM during image-level recovery

Target audience

This guide is intended for users who are responsible for the backup and recovery of VMs. Familiarity with Windows Server administration and the operating system (OS) under which Plug-in *for Hyper-V* is running is assumed. Advanced knowledge of VMs is useful for defining an efficient backup-and-recovery strategy.

Recommended additional reading

Quest recommends that you have the following documentation available for reference while setting up and using this plug-in.

- **Microsoft Hyper-V: Backing Up and Restoring Virtual Machines:** http://msdn.microsoft.com/en-us/library/dd405549(VS.85).aspx

- **NetVault Backup documentation:**

  - *Quest NetVault Backup Installation Guide*: This guide provides details on installing the NetVault Backup Server and Client software.

  - *Quest NetVault Backup Administrator's Guide*: This guide explains how to use NetVault Backup and describes the functionality common to all plug-ins.

  - *Quest NetVault Backup CLI Reference Guide*: This guide provides a description of the command-line utilities.

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

Installing or removing the plug-in

* Deploying the plug-in
* Prerequisites for installing the plug-in
* Prerequisite for using SMB file servers
* Installing the plug-in in a standalone deployment
* Licensing the plug-in
* Removing the plug-in from a standalone deployment

Deploying the plug-in

Deploying the plug-in in standalone, clustered (high-availability), and SMB environments is almost identical because Plug-in *for Hyper-V* is installed on the Hyper-V Server. The Hyper-V Server is the physical machine that hosts the VMs—the plug-in is **not** installed on each of the protected VMs.

Figure 1. Deployment overview

● NetVault Backup Software
◆ Plug-in for Hyper-V for Windows
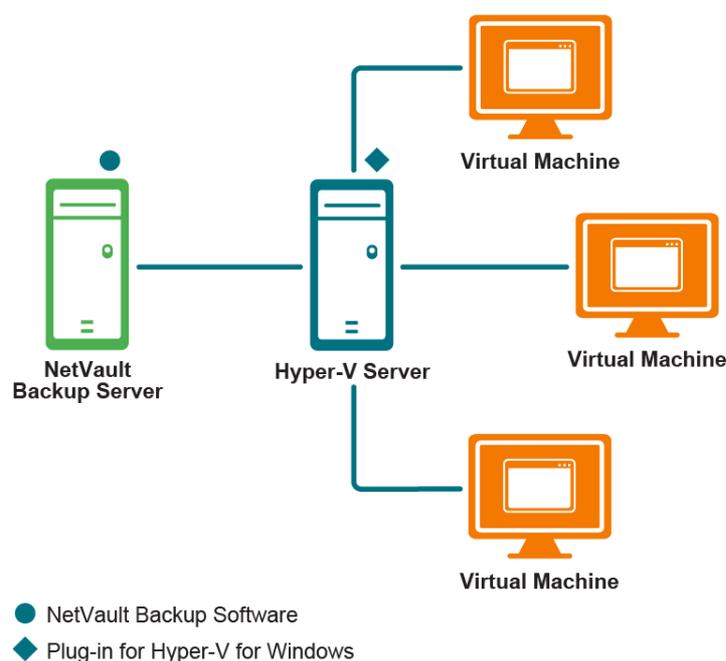
Failover clustering

Using the failover cluster Fully Qualified Domain Name (FQDN) or IP address, the plug-in identifies the current node that is in control of the Hyper-V Cluster and targets it for backup jobs.

Use of a Virtual Client is not plug-in-specific. For complete information, see the *Quest NetVault Backup Administrator's Guide*. Be aware of the following guidelines:

- Always use the information provided in the *Quest NetVault Backup Administrator's Guide* to ensure that management of the plug-in is conducted using the Virtual Client; do *not* use the physical nodes that are members of the Virtual Client to conduct cluster- or client-management activity. Use the physical node to set up backup jobs for local VMs that are not part of the cluster.

- To make the plug-in available for installation to all nodes that are part of a Virtual Client, copy the **".npk"** installation file to the NetVault Backup Server. The path to copy the file is *<NetVaultBackupHome>***\packages\standard**.

- If you later upgrade or update the plug-in, you must:

    - Delete the existing ".npk" file from the same location identified earlier.

    - Copy the ".npk" installation file to the same location identified earlier.

    - Remove the Virtual Client that was created with the previous version of the plug-in.

    - Create a new Virtual Client using the upgraded version of the plug-in. Assign the original Virtual Client's name to the new Virtual Client. If you change the name, you cannot run backup jobs that were defined for the original Virtual Client.

- Do *not* do either of the following:

    - Do not use the **Manage Clients** option to remove the plug-in from clients that are part of a Virtual Client.

    - Do not use the **Manage Clients** option to reinstall the plug-in.

Prerequisites for installing the plug-in

Before installing Plug-in *for Hyper-V*, verify that the following prerequisites are met on the system designated as the Hyper-V Server.

- **Install Windows PowerShell and configure the execution policy:** Before you can use the plug-in, PowerShell 3.0 or later must be installed on the Hyper-V Server, and the execution policy for the local machine must be set to RemoteSigned. By default, Windows Server 2008 R2 includes PowerShell 2.0, which you must update to 3.0 or later. Windows Server 2012, 2012 R2, and 2016 include 3.0, 4.0, and 5.0, respectively.

    After you have installed it, run the following command from the PowerShell command prompt:

    ```
    Set-ExecutionPolicy RemoteSigned
    ```

- **Install Integration Services:** The **Integration Services** must be installed on the VM from the Hyper-V Server. Integration Services are installed by default on Windows-based VMs.

- **Install NetVault Backup Server and Client software:** At a minimum, the NetVault Backup Client software must be installed on the Hyper-V Server. For instructions on installing the NetVault Backup Client software and adding the Hyper-V Server as a Heterogeneous Client, see the *Quest NetVault Backup Installation Guide*.

Prerequisite for using SMB file servers

If the Hyper-V VMs that you back up using the plug-in reside on an SMB file server, verify that the following server roles are enabled. If not, use the Server Manager in Windows Server to add these roles.

Table 1. Required roles for use with VMs stored on SMB file servers

Name	Type	Path
File and Storage Services	Role	File and Storage Services
File and iSCSI Services	Role Service	File and Storage Services\File and iSCSI Services
File Server	Role Service	File and Storage Services\File and iSCSI Services\File Server
File Server VSS Agent Service	Role Service	File and Storage Services\File and iSCSI Services\File Server VSS Agent Service

Installing or upgrading the plug-in

The installation process depends on whether you are deploying the plug-in in a standalone or clustered environment:

- Installing the plug-in in a standalone deployment
- Installing the plug-in in a clustered deployment

Installing the plug-in in a standalone deployment

1 Access the **NetVault Configuration Wizard** or **Manage Clients** page.

> **i** | **NOTE:** If the selected clients are all the same type, you can use the configuration wizard to install the plug-in on multiple clients at the same time. When you select multiple clients, verify that the plug-in binary file is compatible with the OS and platforms of the target clients. From the **Manage Clients** page, you can only select one client for plug-in installation.

- To access the **NetVault Configuration Wizard** page:

    a In the Navigation pane, click **Guided Configuration**.

    b On the **NetVault Configuration Wizard** page, click **Install Plugins**.

    c On the next page, select the applicable clients.

- To access the **Manage Clients** page:

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

    b On the **Manage Clients** page, select the machine that contains the Hyper-V Server, and click **Manage**.

    c On the **View Client** page, click the **Install Plugin** button ( ).

2 Click **Choose Plug-in File**, and 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.

Based on the OS in use, the path for this software may vary on the installation CD.

3 Select the file entitled **"hv-x-x-x-x.npk,"** where **xxxxx** represents the version number and platform, and click **Open**.

4   To begin installation, click **Install Plugin**.

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

Installing the plug-in in a clustered deployment

Installation of the plug-in in a clustered environment is completed through the creation of a **Virtual Client** on the NetVault Backup Server. A Virtual Client is a group of nodes within the cluster. NetVault Backup Server views this group as a *single* client that is created to back up clustered VMs. During the Virtual Client creation process, the plug-in is transferred from the NetVault Backup Server to selected nodes within a cluster and installed there.

Creating a Virtual Client

The Virtual Client creation process is not plug-in-specific, and you can find complete details in the *Quest NetVault Backup Administrator's Guide*. However, consider the following point during the Virtual Client creation process:

Assign a name to the Virtual Client: Quest strongly recommends that you use the cluster network name—that is, the FQDN—assigned to Hyper-V Cluster as the NetVault Backup Virtual Client name. When you browse a Virtual Client, NetVault Backup locates the node currently in control of the clustered application, and reveals the Hyper-V Server—for example, on the **NetVault Backup Selections** page. With a Virtual Client name set up as the Hyper-V Server network name, you can recognize the Hyper-V Server for which the Virtual Client was created.

Licensing the plug-in

The plug-in requires separate license keys for clustered and standalone environments. For a clustered environment, apply the license to the physical client for the NetVault Backup Server; if you also use the same physical client for standalone protection, apply a separate standalone license to the same client. For information about the licensing process, including how to obtain the proper license keys, see the *Quest NetVault Backup Installation Guide*.

Removing the plug-in from a standalone deployment

For details on removing the Plug-in *for Hyper-V* in a clustered environment, see the applicable topic on working with client clusters in the *Quest NetVault Backup Administrator's Guide*.

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

2   On the **Manage Clients** page, select the applicable client, and click **Manage**.

3   In the **Installed Software** table on the **View Client** page, select **Plug-in *for Hyper-V***, and click the **Remove Plugin** button ( ).

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

Configuring the plug-in

- Configuring settings for a clustered deployment
- Enabling communication between VMs in clustered environment

Configuring settings for a clustered deployment

If you are using a standalone deployment or a clustered deployment that is based on a version of Windows Server prior to 2012, the **Configure** option is visible but not accessible in the context menu; no configuration is required.

If you are using a Windows Server 2012, 2012 R2, or 2016 environment, an extra option, **Show VMs being managed by a cluster**, might be visible on the **Plug-in** *for Hyper-V* tab. This option is only available for a physical machine if the host is part of a Hyper-V Cluster.

1 In the Navigation Pane, click **Create Backup Job**, and click **Create New** next to the **Selections** list.

2 In the selection tree, open the applicable client node.

3 Click **Hyper-V Plugin**, and select **Configure** from the context menu.

4 If the **Show VMs being managed by a cluster** option is available, select it *only* if the host is *not* managed by a Virtual Client.

> **i** | **IMPORTANT:** Quest recommends that you use Virtual Clients to manage clustered hosts and that you do *not* enable this option.

5 To configure a Virtual Client, enter the following information:

- **Cluster Domain:** Enter the Windows domain of the cluster.
- **Username:** Enter the account name that has the permission and privileges required for administering the cluster and Hyper-V; that is, a domain-administrator account.
- **Password:** Enter the password associated with the user specified in the preceding field.

6 To save the settings, click **OK**.

Enabling communication between VMs in clustered environment

If you are using NetVault Backup Client software 10.0.5 or earlier, after installing and configuring a NetVault Backup Virtual Client, ensure that the Virtual Client can manage all hosts included in a cluster. Also ensure that each host can access any other host included in the cluster.

> **i** | **TIP:** For clusters with many hosts, you can create a script file that contains all the following commands. You can then run the script on each host.

1   From a command prompt on the first host included in the cluster, type:

    cd <NetVaultBackupInstallDirectory>\util

The default installation directory is **C:\Program Files (x86)\Quest\NetVault Backup\util** or the corresponding directory under **C:\Program Files**.

2   For each host that was added to the Virtual Client, type:

    nvclientaccess -client <clientName> -password <clientPassword>

where *<clientName>* is the name of the next physical client that was added to the Virtual Client and *<clientPassword>* is the same one that you specified when you added the client to the NetVault Backup Server.

For example, if you have three hosts, run nvclientaccess -client *<clientName>* password *<clientPassword>* on Host1 with *<clientName>* as **Host2**. Run nvclientaccess -client *<clientName>* password *<clientPassword>* again on Host1 with *<clientName>* as **Host3**.

3   Repeat these steps for every host included in the cluster.

Backing up data

- Backing up data: an overview
- Adding patterns to use for inclusion and exclusion of VMs
- Performing backups
- Using backup policies

Backing up data: an overview

Plug-in *for Hyper-V* supports image-level backups and restores. An image-level backup includes all disk and configuration files associated with a given VM, enabling restoration of the entire VM. You can use this type of backup if hardware fails or VM disk files are deleted. Using the file-level indexing feature, you can also back up images with file-level indexing enabled and later perform file-level restores. As described in the next topic, there are network-connectivity issues with non-VSS-compliant VMs.

Important considerations

- Do not schedule multiple image-level backups to run simultaneously on the same VM. Running these backup types for the same VMs simultaneously may result in failure of one of these jobs. The plug-in allows only a single snapshot for a VM at any given point.

- Backups cannot be performed if the Hyper-V Server or Cluster is in maintenance mode or inaccessible.

- Plug-in *for Hyper-V* supports the backup of databases that reside on VMs provided the databases are in a consistent state. The vendor, such as the VSS Writer for SQL Server and Oracle, must support the handling of consistent data before it provides the data to the plug-in. If an inconsistent snapshot error—VSS_E_WRITERERROR_INCONSISTENTSNAPSHOT—occurs, the plug-in notes this information in the debug trace and continues with the backup process.

- The plug-in uses installed VSS hardware providers on the backup host. If a provider-related error occurs, the plug-in tries to complete the backup using the applicable software provider.

- Online image-level backups are not fully supported for the following VMs:

  - Non-VSS-compliant VMs, such as Linux
  - Windows Server 2012 VM hosted by Windows Server 2008 R2

  During backup, these VMs are placed in a Saved State, which causes the VM to pause. The VM resumes activity after the snapshot is taken.

- In addition to supporting the use of spaces within the name of a VM, the following characters are also supported:: A–Z, a–z, 0–9, ! # % ^ ( ) - _ + = . { } ;

  Be aware of the following guidelines and exceptions:

  - Names are case-sensitive.
  - The plug-in does not support names that end with a whitespace.
  - The inclusion and exclusion features do not support the use of spaces at the beginning or end of a VM name pattern.

▪ The behavior of unsupported characters is unpredictable.

• If your environment uses a Windows Server 2016 host or cluster, you can complete Full, Incremental, and Differential Backups. A Full Backup performs a complete backup of the set of items that have been selected. An Incremental Backup backs up the set of items that have changed since the last backup. A Differential Backup backs up the items that have changed since the last Full Backup. The plug-in takes advantage of RCT, which tracks the changed blocks.

Adding patterns to use for inclusion and exclusion of VMs

In addition to selecting VMs from the backup selection tree, you can create and store patterns of VM names for inclusion and exclusion from backup jobs. When you specify VM name patterns to include or exclude, or both, the plug-in stores them with a Backup Selection Set. When you submit a backup job, you can select the set with the stored patterns. The plug-in then populates the backup list with VMs whose names match the specified pattern or patterns.

> **i** | **IMPORTANT:** During a backup job, the plug-in gives preference to VMs located under the localhost or
> | <clusterName> node versus information specified using the inclusion and exclusion feature.

To add patterns:

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

2 Next to the **Selections** list, click **Create New**.

3 In the list of plug-ins on the **NetVault Backup Selections** page, open **Hyper-V Plugin**.

Below the **localhost** or ***<clusterName>*** node, the plug-in displays two additional items, **Specify Inclusions** and **Specify Exclusions**. Clicking **Specify Inclusions** displays additional options, such as **Show include patterns** and **Add include pattern**. After you have selected **Show include patterns**, the option changes to **Hide include pattern**. Clicking **Specify Exclusion** displays additional options, such as **Show exclude patterns** and **Add exclude pattern**. After you have selected **Show exclude patterns**, the option changes to **Hide exclude pattern**. As you add patterns, the plug-in displays an informational node for each saved pattern.

4 To add a pattern that you want to include in a backup, click **Add include pattern**.

5 In the **Enter include pattern** dialog box, type the pattern that you want the plug-in to search for.

Use the asterisk (*) wildcard character as part of the pattern to search for. For example, if you want the plug-in to locate all VMs that start with **SQL**, type **SQL***.

> **i** | **NOTE:** If you do not specify an inclusion pattern but you do specify an exclusion pattern, the plug-in
> | automatically assumes that all VMs are included; that is, the plug-in applies an inclusion pattern of *.
> | The plug-in then applies the exclusion pattern to the comprehensive list of included VMs.

6 To save the pattern, click **OK**.

7 Repeat Step 4 through Step 6 for each pattern that you want to add for inclusion.

8 To add a pattern that you want to exclude from a backup, click **Add exclude pattern**.

9 In the **Enter exclude pattern** dialog box, type the pattern that you want the plug-in to search for using the same guidelines described previously for include patterns, and then click **OK**.

Continuing with the previous example, if you want the plug-in to locate all VMs that start with **SQL** but exclude specific versions of the SQL Server VMs, type the complete name of the VM that you want to exclude, such as **SQLQATest1**.

10 Repeat Step 8 and Step 9 for each pattern that you want to add for exclusion.

11 When you are finished, select all required VMs under **localhost** or ***<clusterName>*** node and select the applicable inclusion and exclusion patterns.

12 Click **Save**, enter a name in the **Create New Set** dialog box, and click **Save**.

The name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Windows, there is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

Performing backups

The procedure for performing image-level backups, with or without indexing, includes the steps outlined in the following topics.

- Selecting data for a backup
- Setting backup options
- Finalizing and submitting the job

Selecting data for a backup

You must use sets—Backup Selection Set, Backup Options Set, Schedule Set, Target Set, and Advanced Options Set—to create a backup job.

> **i** | **TIP:** To use an existing set, click **Create Backup Job**, and select the set from the **Selections** list.

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

You can also start the wizard from the Guided Configuration link. In the Navigation pane, click **Guided Configuration**. On the **NetVault Configuration Wizard** page, click **Create backup jobs**.

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

Assign a descriptive name that lets you easily identify the job when monitoring its progress or restoring data. The job name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Windows, there is no length restriction. However, a maximum of 40 characters is recommended.

3 Next to the **Selections** list, click **Create New**.

4 In the list of plug-ins on the **NetVault Backup Selections** page, open **Hyper-V Plugin**, and expand the **localhost** or ***<clusterName>*** node to display the list of VMs.

5 Select the **localhost** or ***<clusterName>*** node to include every VM that exists at the time of backup, or select individual VMs to back up a specific set of VMs.

> **i** | **IMPORTANT:** If VMs are explicitly selected with a green check mark, only the selected VMs are included in the backup, even if new VMs have been added to the Hyper-V Server since the last backup.
>
> If you are performing an Incremental or Differential Backup for a Windows Server 2016 environment, select the Backup Selection Set of the last Full Backup.

6 Click **Save** to save the set, enter a name in the **Create New Set** dialog box, and click **Save**.

The name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Windows, there is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

Setting backup options

The next step involves creating the Backup Options Set or selecting an existing one.

> **i** | **TIP:** To use an existing set, in the **Plugin Options** list, select the set that you want to use.

There are multiple options available: **Backup Type**, **Maximum Number of Parallel Streams**, **Enable indexing for File Level Recovery**, and **Enable Active Block Mapping**.

- The **Backup Type** section is available if the Backup Selection Set is based on a Windows Server 2016 host or cluster. Because the plug-in uses RCT to ensure that only changed blocks are sent to the Incremental or Differential Backup stream, you can indicate whether you want to complete a Full, Incremental, or Differential Backup. A Full Backup performs a complete backup of the set of items that have been selected. A Full Backup does not depend on any other backup and can be restored in a single step. An Incremental Backup backs up the set of items that have changed since the last backup. A Differential Backup backs up the set of items that have changed since the last Full Backup.

- The default for **Maximum Number of Parallel Streams** is **4** for Windows Server 2012, 2012 R2, and 2016, and **1** for Windows Server 2008 R2. If your environment uses Windows Server 2012, 2012 R2, or 2016, you can increase the setting up to 32.

- The **Enable indexing for File Level Recovery** option is disabled by default. Indexing lets you restore individual files and directories from image-level backups of VMs.

  File-level indexing is available to VMs that use the following file systems:

  - Windows**:** NTFS

  - Extended file systems: EXT2, EXT3, and EXT4

  - XFS on Linux and UNIX

  File-level indexing does not affect the size of the backup. However, it increases the size of the backup index and the total backup time. The amount of time taken to perform file-level indexing depends on several factors. These factors include the number of files, fragmentation of files on the volumes, network traffic, and load on the Hyper-V Server.

  The plug-in also supports volumes managed by Logical Volume Manager (LVM) on Linux-based systems and Logical Disk Manager (LDM) on Windows-based systems as single or spanned disks. It does not support Windows Server 2012 or 2016 ReFS (Resilient File System) and Striped Disks or Storage Spaces.

  If you intend to use file-level recovery for Linux- or UNIX-based VMs hosted on Windows Hyper-V Servers, the names of files, directories, and volumes cannot contain keyboard combinations of Alt+0 through Alt+32, nor can they contain the following characters: \\ \\ / [ ] : | < > + ; = . ? " (For more information, see https://technet.microsoft.com/en-us/library/cc956689.aspx.)

- The **Enable Active Block Mapping** option is enabled by default. To reduce network and storage requirements, ABM technology provides filters to remove zeroed blocks during backups. Removing these blocks reduces the backup size and the amount of data transferred over the network. If you disable ABM, the plug-in automatically uses zero-block exclusion to ensure that the zeroed blocks are excluded during backups.

  ABM is available to VMs that use the following file systems:

  - Windows**:** NTFS

  - Extended file systems: EXT2, EXT3, and EXT4

To select the backup options:

1. Next to the **Plugin Options** list, click **Create New**.

2. If your environment uses Windows Server 2016 or later, in the **Backup Type** section, select **Full**, **Incremental**, or **Differential**, whichever is applicable.

3. In the **Maximum Number of Parallel Streams** box, enter the applicable number.

If your environment uses Windows 2012, 2012 R2, or 2016, you can enter up to 32 to improve load balancing or performance, or both. However, if you are using tape drives, the number of parallel backup streams selected *must* be equal to or less than the number of available drives.

4   If you might want to perform restores at the file level, select the **Enable indexing for File Level Recovery** option.

5   If you want to disable ABM for any reason, clear the **Enable Active Block Mapping** option.

6   Click **Save**, specify a name for the set in the **Create New Set** dialog box, and click **Save**.

The name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Windows, there is no length restriction; however, a maximum of 40 characters is recommended.

Finalizing and submitting the job

1   Use the **Schedule**, **Target Storage**, and **Advanced Options** lists to configure any additional required options.

2   Click **Save** or **Save & Submit**, whichever is applicable.

> **i** | **TIP:** To run a job that you have already created and saved, select **Manage Job Definitions** in the Navigation pane, select the applicable job, and click **Run Now**.

You can monitor progress on the **Job Status** page and view the logs on the **View Logs** page. For more information, see the *Quest NetVault Backup Administrator's Guide*.

> **i** | **NOTE:** If you add a VM to the Backup Selection Set after the initial Full Backup or any subsequent Incremental or Differential Backup, the plug-in does the following:
> - Creates a Full Backup of the new VM if you run an Incremental or Differential Backup.
> - For subsequent Differential Backups, runs a Full Backup of the new VM if RCT is supported for the VM disks.
> - For subsequent Incremental Backups, runs an Incremental Backup of the new VM if RCT is supported for the VM disks.
>
> Example:
> 1   Select any VM, for example, MyVM1, and create a Backup Selection Set, for example, MySelectionSet.
> 2   Run a Full Backup using MySelectionSet.
> 3   Modify MySelectionSet, and add a VM, for example, MyVM2.
> 4   Run an Incremental or Differential Backup using MySelectionSet.
> 5   Run another Incremental Backup or Differential using MySelectionSet.
>
> As a result, when you select **Incremental Backup** for the **Backup Type**, the plug-in creates a Full Backup of MyVM2 in Step 4 and an Incremental of that VM in Step 5. When you select Differential Backup for the Backup Type, the plug-in creates a Full Backup of MyVM2 in Step 4 and Step 5.

Using backup policies

You can use policies to submit one or more jobs that target one or more similar clients. The following procedure outlines the basic process. For more information, see the *Quest NetVault Backup Administrator's Guide*.

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

2   On the **Manage Policies** page, click the browse button, and select **Add**.

3   Enter a **Policy Name**, and click **Add job**.

4   On the **Create Policy Job** page, complete the applicable sections, and click **Save**.

> **i** | **TIP:** You can use the inclusion- and exclusion-pattern feature to select which VMs to back up. For more information, see Adding patterns to use for inclusion and exclusion of VMs.

5 Click **Add Clients**, and select the applicable Hyper-V clients.

6 To save the policy definition, click **Save Policy**.

Restoring data

- Restoring image-level backups
- Performing file-level recovery
- Using other restore procedures

Restoring image-level backups

You can perform an image-level recovery, which restores the entire VM and establishes a new version of that VM in Hyper-V Manager. The process restores files either to their original location or to a location specified by you when you create a selection set and enter information in the Rename/Relocate dialog box. For instructions on performing restores from a previous version of the plug-in, see the user's guide for the corresponding version.

> **i** | **IMPORTANT:** The plug-in flattens Hyper-V checkpoints when a backup is taken and later retrieved using an image-level recovery process. For example, if you back up a base disk that has two checkpoint files that the VM is running, the VM is restored with only one base disk; however, all content from the checkpoint files is included.

This process includes the steps outlined in the following topics:

- Prerequisites for image-level restore
- Selecting data for an image-level restore
- Setting image-level restore options
- Finalizing and submitting an image-level restore job

Prerequisites for image-level restore

Verify that the following prerequisites are met:

- If you do not intend to use the **Rename** feature, Quest recommends that you delete the original VM *before* you perform recovery.

- If you do not intend to use the **Rename** feature and you do not intend to delete the original VM, ensure that the VM is turned off. Also ensure that none of its files are open *before* you perform recovery. The plug-in's recovery process overwrites the existing VM, and write-access is required.

- Before recovering a VM to different host, verify that the VM resources in the Configuration File match the resources available on the host. In particular, verify that file locations and network-adapter names match. If they do not match, the VM might not get turned on. If resources are different on the target host, complete the recovery manually.

Selecting data for an image-level restore

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

2   On the **Create Restore Job—Choose Saveset** page, select **Plug-in** *for Hyper-V* from the **Plugin Type** list.

3   To filter the items displayed in the saveset table further, use the **Client**, **Date**, and **Job ID** lists.

    The table displays the saveset name (job title and saveset ID), creation date and time, and size. By default, the list is sorted by creation date.

4   In the saveset table, select the applicable item.

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

5   Click **Next**.

6   On the **Create Selection Set** page, select the **localhost** or ***<clusterName>*** node to display the VMs, and then select one or more VMs to restore.

    > **i** | **NOTE:** You cannot select individual files under a VM node. If you enabled indexing during backup and want to restore individual files, use file-level recovery. For more information, see Performing file-level recovery.
    >
    > If you select an Incremental or Differential Backup, the plug-in automatically restores all savesets in the backup sequence up to the selected saveset.

Setting image-level restore options

If you are restoring a backup that was created using an earlier version of the plug-in, complete the steps outlined in the user's guide for the corresponding version; this section only applies to the current version of the plug-in.

1   On the **Create Selection Set** page, click **Edit Plugin Options.**

2   Select one of the following parameters:

    ▪   To restore the entire VM and establish a new version of that VM in Hyper-V Manager, select **Recover entire VM to host and/or cluster**.

    ▪   To restore specific files to a location specified by you, select **Restore files from inside VM guest OS to specified location**.

> **!** | **CAUTION: The Delete working directory contents after recovery option is selected by default. Do** *not* **clear this option unless instructed to do so by Quest support.**

3   If you selected the **Recover entire VM to host and/or cluster** option and you want the recovered VM to be restarted automatically, select the **Power-on VM after recovery** option.

    > **i** | **NOTE:** The **Power-on VM after recovery** option is only available if you select a backup saveset that was created using version 11.4 or later of the plug-in.

Finalizing and submitting an image-level restore job

The final steps include setting additional options on the Schedule, Source Options, and Advanced Options pages, submitting the job, and monitoring the progress through the Job Status and View Logs pages. These pages and options are common to all NetVault Backup Plug-ins. For more information, see the *Quest NetVault Backup Administrator's Guide*.

1   To save the settings, click **Ok**, and then click **Next**.

2   In **Job Name**, specify a name for the job if you do not want to use the default setting.

   Specify a detailed title that lets you easily identify the job when monitoring its progress. The job name can contain alphanumeric and nonalphanumeric 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 **Target Client** list, select the machine on which you want to restore the data.

   If you select a Virtual Client, the plug-in restores the VM to the host that is managing the cluster.

   > **i** | **TIP:** You can also click **Choose**, and then locate and select the applicable client in the **Choose the Target Client** dialog box.

4   Use the **Schedule**, **Source Options**, and **Advanced Options** lists to configure any additional required options.

5   Click **Save** or **Save & Submit**, whichever is applicable.

   You can monitor progress on the **Job Status** page and view the logs on the **View Logs** page. For more information, see the *Quest NetVault Backup Administrator's Guide*.

   > **i** | **NOTE:** If you restore a saveset that was created with Plug-in *for Hyper-V* version 10.5 or earlier, the Virtual Hard Disk (VHD or VHDX) files are stored in *<installationDirectory>*\tmp\Hyper-V.

Performing file-level recovery

If you select the **Enable indexing for File Level Recovery** option during backups generated using version 10.5 or later of the plug-in, you can restore individual files and directories from an image-level backup. The process includes the steps outlined in the following topics:

- Selecting data for a file-level restore
- Setting file-level restore options
- Finalizing and submitting a file-level restore job

> **i** | **IMPORTANT:** The plug-in does not support file-level restores of compressed or sparse Windows Server NTFS files, Linux sparse files, or files that reside on VMs using Storage Spaces.
>
> You can only select one VM at a time for a file-level restore job. To recover files from more VMs, run more file-level restore jobs.

Selecting data for a file-level restore

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

2   On the **Create Restore Job—Choose Saveset** page, select **Plug-in *for Hyper-V*** from the **Plugin Type** list.

3   To filter the items displayed in the saveset table further, use the **Client**, **Date**, and **Job ID** lists.

    The table displays the saveset name (job title and saveset ID), creation date and time, and size. By default, the list is sorted by creation date.

4   In the saveset table, select the applicable item.

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

> **i** | **IMPORTANT:** If your restore job includes a Linux VM, the **Enable indexing for File Level Recovery** option was selected during backup, the VM contains two files in the same directory that use the same name—differentiated only by capitalization—and you must restore both files, create and run *two* restore jobs. In one job, include all applicable files and one of the files with the same name. In the other job, include only the other file that uses the same name, and specify a different location for the restore. Because Windows is not case-sensitive, it does not recognize the distinction between the two files; this workaround prevents the restore from completing with warnings and ensures that both versions of the file are recovered.

5   Click **Next**.

6   On the **Create Selection Set** page, expand the **Volumes** node to display the indexed partitions, directories, and files contained within, and then select one or more indexed items for the restore.

> **i** | **IMPORTANT:** Select only the applicable folders or files, or both, located under the Volumes node. Do not select the VM node itself.
>
> Do *not* use the **Rename** option in the context menu. Although the option is displayed, the plug-in ignores the **Rename** field and the restore completes with warnings.

Setting file-level restore options

1   On the **Create Selection Set** page, click **Edit Plugin Options**, and select the **Restore files from inside VM guest OS to specified location** option.

2   In the **Target directory** box, enter the full path to where you want the files restored, and verify that the path exists before continuing.

    For a standalone configuration, the plug-in recovers files to the Hyper-V host. For a clustered configuration, the plug-in recovers files to the host (Virtual Client) that is managing the cluster. In both cases, the plug-in defaults to placing all the files in **C:*<directoryNamedAfterVM>***.

3   To save the settings, click **Ok**, and then click **Next**.

Finalizing and submitting a file-level restore job

The final steps include setting additional options on the Schedule, Source Options, and Advanced Options pages, submitting the job, and monitoring the progress through the Job Status and View Logs pages. These pages and options are common to all NetVault Backup Plug-ins. For more information, see the *Quest NetVault Backup Administrator's Guide*.

1 To save the settings, click **Ok**, and then click **Next**.

2 In **Job Name**, specify a name for the job if you do not want to use the default setting.

Specify a detailed title that lets you easily identify the job when monitoring its progress. The job name can contain alphanumeric and nonalphanumeric 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 Use the **Schedule**, **Source Options**, and **Advanced Options** lists to configure any additional required options.

4 Click **Save** or **Save & Submit**, whichever is applicable.

You can monitor progress on the **Job Status** page and view the logs on the **View Logs** page. For more information, see the *Quest NetVault Backup Administrator's Guide*.

Using other restore procedures

This topic describes other restore operations that you can perform with the plug-in:

- Searching for files in savesets

- Renaming and relocating a VM during image-level restore

- Restoring data to an alternate Hyper-V Server

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 Backup 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 regular expressions instead of Elasticsearch in the **Search String** box, select this check box.

- **Use legacy search method**: If both cataloged and non-cataloged savesets are included in the search, the plug-in displays this check box.

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

If only cataloged savesets are included in the search or if **Use legacy search method** is cleared, the catalog search is used.

3   To search in one or more specific savesets, select the applicable savesets, and click **Search**.

If you do not select a saveset, all savesets are included in the search. On the **Search Results** page, you can view and filter the savesets that contain the specified files or data items.

4   Select the items you want to restore.

You can only restore items from one saveset.

5   Click **Restore selected items**.

6   Complete the additional steps outlined in Restoring image-level backups or Performing file-level recovery, whichever is applicable.

Renaming and relocating a VM during image-level restore

The plug-in lets you restore a VM to a different name and location. The renaming of a VM can be useful if you do not want to overwrite the existing version.

> **i**   **IMPORTANT:** If your environment uses Windows Server 2008 R2 and you use the Rename feature, the recovery does *not* include network adapters that the VM backup included.

1   In the Navigation pane, click **Create Restore Job**, select **Plug-in** *for Hyper-V* from the **Plugin Type** list, select the applicable saveset, and click **Next**.

2   On the **Create Selection Set** page, select the **localhost** or ***<clusterName>*** node to display the VMs, and then select the VM to restore.

3   With the VM selected on the **Create Selection Set** page, select **Rename** from the context menu.

4   In the **Rename/Relocate** dialog box, complete both fields (required):

  ▪   **Rename:** Enter the new name.

  ▪   **Relocate:** Enter the new full path. If you are recovering clustered VMs, verify that the path is part of the cluster.

5   To apply your changes, click **Ok**.

The VM name is accompanied by the rename and relocate information in parentheses.

6   Continue with the restore procedure as explained in Restoring image-level backups.

7   If your environment uses Windows Server 2008 R2 and the original VM included network adapters, use Hyper-V Manager to add the network adapters.

Restoring data to an alternate Hyper-V Server

This functionality is useful during *disaster recovery* operations. You can restore a VM to a different Hyper-V Server by using the steps provided in the following topic.

1   Verify that the following requirements are met on the new restore target before continuing this process.

  ▪   **Same version of Hyper-V Server installed:** The Hyper-V Server version must be the same as the version running on the existing server.

  ▪   **NetVault Backup software and the Plug-in** *for Hyper-V* **installed:** The same version of NetVault Backup software and the plug-in must be installed and configured on the new restore target.

  ▪   **Client added to the NetVault Backup Server:** The target machine must be added to the NetVault Backup Server by using the **Manage Clients** page of the NetVault Backup WebUI.

2   In the Navigation pane, click **Create Restore Job**, select **Plug-in *for Hyper-V*** from the **Plugin Type** list, select the applicable saveset, and click **Next**.

3   Select the VM to be restored from the backup saveset.

4   On the **Create Selection Set** page, click **Edit Plugin Options**, and select the **Recover entire VM to host and/or cluster** option.

    For information about other options available, see Setting image-level restore options.

5   To save the settings, click **Ok**, and then click **Next**.

6   In **Job Name**, specify a name for the job if you do not want to use the default setting.

    Assign a descriptive name that lets you easily identify the job when monitoring its progress. The job name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Windows, there is no length restriction; however, a maximum of 40 characters is recommended.

7   In the **Target Client** list, select the alternate Hyper-V Server.

8   Use the **Schedule**, **Source Options**, and **Advanced Options** lists to configure any additional required options.

9   Click **Save** or **Save & Submit**, whichever is applicable.

    You can monitor progress on the **Job Status** page and view the logs on the **View Logs** page. For more information, see the *Quest NetVault Backup Administrator's Guide*.

10 After the restore is finished, use Hyper-V Manager to import the new VM.

Troubleshooting

- NetVault Backup 10.x service, netvault-pgsql, does not start on Windows
- Backup job fails while using VSS snapshot

NetVault Backup 10.x service, netvault-pgsql, does not start on Windows

To see if the following message is displayed, check the Windows Event Viewer: PDT FATAL: lock file "postmaster.pid" already exists.

NetVault Backup 10.x uses a PostgreSQL database. If the PostgreSQL database does not start, NetVault Backup cannot start. To correct this issue, delete the **"postmaster.pid"** file from the location referenced in the log and restart the NetVault Backup Server. For more information, see https://support.questcom/netvault-backup/kb/122475.

Backup job fails while using VSS snapshot

If a backup job is stopped or fails while a VSS snapshot is being used, the plug-in might not clean up the VSS snapshot. If this issue occurs, you can use a tool such as Microsoft CLI **diskshadow.exe** to remove the snapshot.

1　After DiskShadow is run from a command line with elevated permissions, run the following command at the **DISKSHADOW>** prompt to find and remove the orphaned shadow copy.

```
list shadows all
```

The preceding command lists the current shadow copies.

2　Look through the list of shadow copies noting the **Exposed Locally as** field.

3　Find the field with the exposed directory name that matches the VM name used during the stopped or failed job.

4　Using the **Shadow copy ID** field from that entry, run the following command:

```
delete shadows ID <ShadowCopyIDOfAbortedOrFailedJob>
```

5　Navigate to *<installationDirectory>*\tmp\Hyper-V, and delete the directory with the VM name used during the stopped or failed job.

6　Exit DiskShadow and return to a standard command prompt.

7　To verify that all VSS Writers, especially Hyper-V VSS Writer, are in a Stable State, run the following command:

```
vssadmin list writers
```

The preceding command lists the VSS Writers and their State.

8   If a VSS Writer is listed in a State other than Stable, use Hyper-V Manager to select the host node, and stop and start the **Virtual Machine Manager** service.

About us

We are more than just a name

We are on a quest to make your information technology work harder for you. That is why we build community-driven software solutions that help you spend less time on IT administration and more time on business innovation. We help you modernize your data center, get you to the cloud quicker and provide the expertise, security and accessibility you need to grow your data-driven business. Combined with Quest's invitation to the global community to be a part of its innovation, and our firm commitment to ensuring customer satisfaction, we continue to deliver solutions that have a real impact on our customers today and leave a legacy we are proud of. We are challenging the status quo by transforming into a new software company. And as your partner, we work tirelessly to make sure your information technology is designed for you and by you. This is our mission, and we are in this together. Welcome to a new Quest. You are invited to Join the Innovation™.

Our brand, our vision. Together.

Our logo reflects our story: innovation, community and support. An important part of this story begins with the letter Q. It is a perfect circle, representing our commitment to technological precision and strength. The space in the Q itself symbolizes our need to add the missing piece—you—to the community, to the new Quest.

Contacting Quest

For sales or other inquiries, 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.