


**Dell™ NetVault™ Backup
Plug-in for VMware® 11.2**
User's Guide





© 2016 Dell Inc. All rights reserved.

This product is protected by U.S. and international copyright and intellectual property laws. Dell, the Dell logo, and NetVault are trademarks of Dell Inc. in the United States and/or other jurisdictions. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Red Hat, and Red Hat Enterprise Linux are trademarks or registered trademarks of Red Hat, Inc. in the U.S. and other countries. SLES and SUSE are registered trademarks of SUSE LLC in the United States and other countries. SQL Server, SharePoint, and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. VMware, VMware vCenter, VMware ESXi, VMware vSphere, VMware vSphere Storage vMotion, VMware vSphere Thin Provisioning, and VMware vSphere Storage DRS are registered trademarks or trademarks of VMware, Inc. in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Legend

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

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

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

Plug-in for VMware User's Guide
Updated - November 2016
Software Version - 11.2
VMW-101-11.2-EN-01

Contents

Introduction	6
About Dell™ NetVault™ Backup Plug-in for VMware®	6
Key benefits	6
Feature summary	7
Target audience	7
Recommended additional reading	8
Installing the plug-in	9
About plug-in editions	9
Deployment overview	9
Deploying the plug-in on a physical machine	10
Deploying the plug-in in a virtual machine	10
Prerequisites	11
Installing the plug-in	12
Installing the plug-in using the push installation method (Windows only)	12
Installing the plug-in using the configuration wizard	13
Installing the plug-in from the Manage Clients page	13
Removing the plug-in	13
Configuring the plug-in	14
Adding a server	14
Reconfiguring a server	15
Removing a server	16
Supported transport modes	16
Configuring default settings	17
Enabling or disabling CBT on virtual machines	19
About CBT	19
Enabling CBT on individual virtual machines	20
Disabling CBT on a virtual machine	21
Manually disabling CBT on a virtual machine	21
Enabling or disabling quiescing of virtual machines	21
About quiescing virtual machines	21
Enabling quiescing of virtual machines	22
Disabling quiescing of virtual machines	23
Removing backup preferences for a virtual machine	23
Unlocking a virtual machine	23
Defining a backup strategy	24
About backup methods and types	24
Image-level backups	24
File-level backups	25
Data backed up and restored for different disk types	26
RDM disks	27

Backup and recovery strategy	27
Using the image-level backup method	29
Performing image-level backups	29
Additional notes	35
Backup selection tree icons	37
Restarting a backup job	38
Resetting CBT on a virtual machine	38
Monitoring job progress	39
Using the file-level backup method	40
Performing file-level backups	40
Backup selection tree icons	42
Removing the snapshot and mount folder manually	43
Restoring image-level backups	45
About restoring image-level backups	45
Restoring a full virtual machine or individual virtual disks	46
Prerequisites	46
Restoring data	46
Starting a virtual machine	50
Relocating a virtual machine to an alternate ESXi Server	50
Renaming a virtual machine during restore	51
Performing file-level restores from image-level backups	52
Restoring virtual machine disk and configuration files	54
Restoring data	54
Recovering a virtual machine from the restored files	57
Searching for files in savesets	58
Viewing media list	59
Restoring file-level backups	60
About restoring file-level backups	60
Restoring file-level backups using shared network drives	60
Setting up a network share	60
Restoring data	61
Restoring file-level backups using the native Plug-in for FileSystem	63
Prerequisites	63
Restoring data	64
Searching for files in savesets	65
Viewing media list	65
Troubleshooting	67
Common errors	67
Diagnosing virtual machine issues	73
SOAP messages	73
Generating VDDK logs	74
About Dell	75

Contacting Dell75
Technical support resources75

Introduction

- [About Dell™ NetVault™ Backup Plug-in for VMware®](#)
- [Target audience](#)
- [Recommended additional reading](#)

About Dell™ NetVault™ Backup Plug-in for VMware®

Dell NetVault Backup Plug-in for VMware (Plug-in for VMware) protects your virtual machines from disaster, media failure, and data corruption. Its intuitive, easy-to-use interface gives you a single, centralized console to configure backup and restore policies for your virtual environments. The Plug-in for VMware is integrated with VMware vSphere Storage APIs - Data Protection, formerly known as VMware vStorage APIs for Data Protection (VADP). It allows you to protect your virtual environment without the need for complex scripting or deep understanding of VMware® ESXi™ or VADP internals. The plug-in minimizes downtime by allowing you to restore complete virtual machine images or individual files quickly and reliably with minimal interaction. Through automatic integration with a wide range of backup devices, your virtualized data can be protected and safely stored off-site to meet your disaster recovery and business continuity goals.

① **NOTE:** The NetVault Backup Plug-ins for applications such as SQL Server®, Exchange, or SharePoint® can be combined with the Plug-in for VMware to provide comprehensive protection to your data. The Plug-in for VMware provides bare metal recovery for virtual machines. The application plug-ins simplify data protection of business-critical applications by automating the backup and recovery processes and performing important maintenance functions such as transaction log truncation.

Key benefits

- **Improve VMware® ESXi™ performance with vStorage APIs for Data Protection (VADP):** The Plug-in for VMware works with VADP to provide centralized backups that reduce the load on ESXi Servers and backup traffic on the Local Area Network. This approach is especially important if you need to consolidate hardware. You have the flexibility to protect all your virtual machines deployed on a single ESXi Server or protect multiple ESXi Servers from a single NetVault Backup Client. With VMware vSphere® Storage vMotion® (Storage vMotion), the plug-in is able to protect your virtual machines without requiring detailed knowledge of their location.
- **Increase confidence while deploying VMware:** The Plug-in for VMware protects individual ESXi Servers and a full VMware® vCenter™ environment. It empowers you to create comprehensive, flexible backup policies without the need for complex scripting or deep understanding of VADP. The plug-in provides peace of mind because you simply point-and-click to back up or restore virtual machines. It increases efficiency by consolidating virtual machine backup management into a single environment, improving storage administration efficiency. The solution enhances protection for Windows®-based virtual machines by performing VSS-based snapshots to quiesce applications before backup.

- **Automatic integration of backup devices ensures business continuity:** The Plug-in for VMware supports a wide range of backup devices, allowing you to store your backup data on disk-based storage devices, virtual tape libraries, or physical tape libraries. It ensures that your virtualized environments are protected and stored off-site for disaster recovery purposes. The plug-in also frees administrators from being available 24x7 by allowing less experienced VMware administrators to initiate restores correctly and as quickly as possible, reducing downtime and improving business continuity.
- **Shorten backup windows and improve device utilization:** The Plug-in for VMware speeds virtual machine image-based backups with high-performance Incremental and Differential Backups for data stored in virtual machines. The plug-in integrates with the Changed Block Tracking (CBT) feature to ensure that only blocks that have changed since the initial Full, and the last Incremental or Differential Backup are sent to the current Incremental or Differential backup stream. This integration provides more efficient backups and reduced network demands.

The Plug-in for VMware also preserves the use of the VMware vSphere® Thin Provisioning feature during virtual machine recovery operations to control storage costs and protect data. It performs this task by ensuring virtual machine disks only use the amount of storage they currently need while being able to dynamically add more allocated space as required.

Feature summary

- Supports VADP-based online backups
- Creates VSS-based snapshots for application-consistent backups
- Allows CBT-based Full, Incremental, and Differential image-level backups of virtual machines
- Supports backup and restore of Thin Provisioned Disks
- Enables backup and restore of virtual machines within a Virtual Application (vApp)
- Supports Storage vMotion using the vSphere virtual machine locking API
- Supports backup and restore of virtual machines that are protected using VMware vSphere® Fault Tolerance (vSphere FT), when managed under a vCenter Server by VMware® ESXi™ hosts that use version 6.0 build number 4192238 or later
- Reduces backup windows by backing up multiple virtual machines in parallel
- Reduces network and storage requirements with Active Block Mapping (ABM)
- Allows backups and restores of individual virtual disks
- Allows Full, Differential, and Incremental file-level backups of Windows®-based virtual machines
- Offers ease of use with point-and-click GUI
- Restores complete virtual machine images or individual files
- Relocates virtual machines to alternate VMware® ESXi™ Server
- Allows virtual machine renaming during restore
- Allows file-level restores from image-level backups
- Relocates individual files to alternate directory or virtual machine

Target audience

This guide is intended for users who are responsible for the backup and recovery of VMware virtual machines. Familiarity with VMware® vCenter™ and VMware® ESXi™ Server administration and the operating system under which the virtual machines are running is assumed. Advanced knowledge of VMware would be useful for defining an efficient backup and recovery strategy for the virtual machines.

Recommended additional reading

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

You can download the complete set of NetVault Backup documentation from <https://support.software.dell.com/>.

- **VMware documentation:** You can download the complete set of VMware documentation from <http://www.vmware.com/support/pubs>.

Installing the plug-in

- [About plug-in editions](#)
- [Deployment overview](#)
- [Prerequisites](#)
- [Installing the plug-in](#)
- [Removing the plug-in](#)

About plug-in editions

The Plug-in for VMware is available in two editions:

- **ESXi Server Edition:** The ESXi Server Edition allows you to add one or more standalone VMware® ESXi™ Servers to protect all virtual machines hosted by these servers. You can deploy this edition on any Windows®- or Linux®-based NetVault Backup Client. The number of ESXi Servers that you can add to the client depends on the capacity specified in the purchased license.
- **Enterprise Edition:** The Enterprise Edition allows you to add multiple VMware® ESXi™ or VMware® vCenter™ Servers to protect all virtual machines hosted or managed by these servers. You can deploy this edition on any Windows- or Linux®-based NetVault Backup Client. You can add any number of ESXi or vCenter Servers to a client that runs the Enterprise Edition of the plug-in.

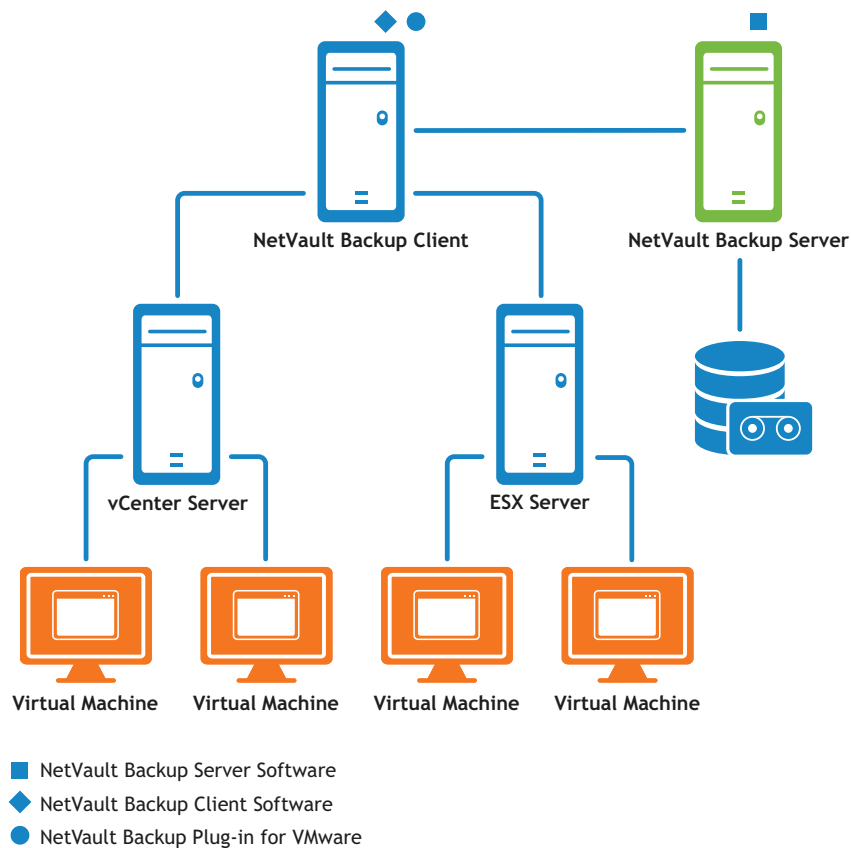
Deployment overview

You can deploy the Plug-in for VMware on any Windows®- or Linux®-based NetVault Backup Client. This client acts as the Backup Proxy. You can select either a physical machine or a virtual machine as the backup proxy.

NOTE: Installing the Plug-in for VMware on a vCenter Server is not recommended.

When you run the plug-in on a vCenter Server, the increased resource demands on the server can affect the overall performance of the server. Also, trace files for large jobs can consume a large amount of disk space.

Figure 1. Plug-in for VMware deployment



Deploying the plug-in on a physical machine

If the virtual machine disks are stored on Fibre Channel SAN, iSCSI SAN, or Serial Attached SCSI (SAS) storage devices, you can select a physical machine to host the plug-in. In this type of deployment, backups are offloaded to the physical backup proxy.

This type of deployment supports SAN and LAN (network block device [NBD] or NBDSSL [encrypted]) transport modes. For more information about transport modes, see [Supported transport modes](#).

Deploying the plug-in in a virtual machine

You can also configure one of the virtual machines as the backup proxy to host the plug-in. This configuration works with all types of storage devices (including local storage on the VMware® ESXi™ Server).

This type of deployment supports HotAdd and LAN (NBD or NBDSSL) transport modes. For more information about transport modes, see [Supported transport modes](#).

① | **NOTE:** If VMware vSphere® Storage DRS™ is enabled for your datastore cluster, exclude the backup proxy from the Storage DRS before you run a job using the HotAdd mode.

Prerequisites

Before installing the plug-in, verify that the following requirements are met:

- **Install NetVault Backup Client:** Install the NetVault Backup Client software on the physical or virtual machine on which you want to install the Plug-in for VMware. For LAN-free backups, you must obtain a NetVault Backup SmartClient license. For more information about installing the NetVault Backup Client, see the *Dell NetVault Backup Installation Guide*.

① **NOTE:** To perform file-level backups of virtual machines, you must install the plug-in on a Windows®-based client. The plug-in version for the Linux® OS does not support this feature.

- **Install applicable libraries on Linux-based clients:** On Linux-based clients, the Plug-in for VMware and Virtual Disk Development Kit (VDDK) require the following libraries:

- **Plug-in for VMware requirements:** On Linux-based systems, the Plug-in for VMware requires the following libraries:

- libssl.so
- libcrypto.so

- **VDDK requirements:** On Linux-based systems, VDDK requires the following libraries:

- **libexpat.so.1.5.2:** To use advanced transport modes, install the XML library **libexpat 1.95.8** on the client. Without the correct version of this library, the jobs that use advanced transport modes may fail.

During installation, the plug-in automatically creates a symbolic link (**libexpat.so.0**) to the library if this link is not present on the client. The plug-in only checks the `/lib64` and `/usr/lib64` directories where this library is typically installed. If the library is installed in a different directory, you must manually create this link.

To create the symbolic link, type:

```
ln -s libexpat.so.1.5.2 <full file path to the library>
```

For example, if **libexpat.so.1.5.2** is installed in the directory `/lib64`, type:

```
ln -s libexpat.so.1.5.2 /lib64/libexpat.so.0
```

If you encounter any issues with the use of advanced transport modes, verify that the correct version of the **libexpat** package is installed. If necessary, manually create the symbolic link. If binary compatibility is present, you can also use a higher version of the library.

- **C++ Library Package:** Verify that the C++ Library is available on the Linux client. Some old Linux distributions might not include this package in a base installation. On such systems, install the C++ Library Package from the distribution ISO.

- **Additional library requirements:** VDDK also requires the following libraries:

- libgmodule-2.0.so
- libxml2.so
- libgcc_s.so

- **Add the NetVault Backup Client:** Add the designated client to the NetVault Backup Server. For more information about adding clients, see the *Dell NetVault Backup Administrator's Guide*.

Installing the plug-in

You can install the plug-in simultaneously on multiple machines by using the configuration wizard. Alternatively, you can install a plug-in on a single client from the **Manage Clients** page.

The following sections describe the different procedures that you can use to install the plug-in:

- [Installing the plug-in using the push installation method \(Windows only\)](#)
- [Installing the plug-in using the configuration wizard](#)
- [Installing the plug-in from the Manage Clients page](#)

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

On Windows®-based machines, you can use the push installation method to install the plug-in on multiple clients at the same time. You can perform push installations from the NetVault Backup WebUI.

NOTE: To use this method, the NetVault Backup Server must be running on a Windows-based machine. The server version must be 11.0 or later.

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

- **Copy the packages to a shared location:** Copy the plug-in packages to a shared location. Only CIFS shares are currently supported as package stores. The path must be accessible to the NetVault Backup 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 Backup:** After copying the installation packages, configure the shared location details in NetVault Backup. For more information, see the *Dell NetVault Backup 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**, and then complete the following steps:
 - a 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**.
 - b Click **Next**.
- 5 On the **Machines to Have NetVault Software Installed** page, click **Choose Machines**, and select **From available machines**.
- 6 On the **NetVault Machine Details** tab, select the client that you want to add, and click **OK**.
- 7 To add more machines, repeat [Step 5](#) and [Step 6](#).
- 8 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 *Dell NetVault Backup Administrator's Guide*.

Installing the plug-in using the configuration wizard

On Linux[®]-based machines, you can use the configuration wizard to install the plug-in on multiple clients at the same time. NetVault Backup 10.x also supports this method on Windows[®]-based machines.

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 Backup Clients** table, select the clients on which you want to install the plug-in.
- 3 Click **Choose Plug-in File**, navigate to the location of the “.npk” installation file for the plug-in, for example, on the installation CD or the directory to which the file was downloaded from the website.
Based on the operating system (OS) in use, the path for this software may vary on the installation CD.
- 4 Select the file entitled “vmw-w.x.y.z-<platform>.npk,” where w.x represents the version number, y represents the patch level, and z represents the build number, and click **Next**.

After the plug-in is successfully installed, 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.

- 1 In the Navigation pane, click **Manage Clients**.
- 2 In the **NetVault Backup Clients** table, select the client, and click **Manage**.
- 3 In the lower-right corner of the **Installed Plug-ins** table, click the **Install Plugin** button (⊕).
- 4 Click **Choose Plug-in File**, navigate to the location of the “.npk” installation file for the plug-in, for example, on the installation CD or the directory to which the file was downloaded from the website.
Based on the operating system (OS) in use, the path for this software may vary on the installation CD.
- 5 Select the file entitled “vmw-w.x.y.z-<platform>.npk,” where w.x represents the version number, y represents the patch level, and z represents the build number, and click **Install Plugin**.

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

Removing the plug-in

- 1 In the Navigation pane, click **Manage Clients**.
- 2 In the **NetVault Backup Clients** list, select the client, and click **Manage**.
- 3 In the **Installed Plug-ins** table, select **VMware Plugin**, and click the **Remove Plugin** button (⊖).
- 4 In the **Confirm** dialog box, click **Remove**.

Configuring the plug-in

- Adding a server
- Reconfiguring a server
- Removing a server
- Supported transport modes
- Configuring default settings
- Enabling or disabling CBT on virtual machines
- Enabling or disabling quiescing of virtual machines
- Unlocking a virtual machine

Adding a server

Before you can use the plug-in for backups and restores, you must add the applicable VMware® ESXi™ and VMware® vCenter™ Servers to the plug-in, and configure a user account to log on to the server.

- **Standalone ESXi Server setup:** In a VMware environment that consists of one or more ESXi Servers, add the individual hosts to the plug-in, and configure a user account for each server.
- **Cluster setup:** In a cluster setup managed by a vCenter Server, add the vCenter Server to the plug-in, and configure a user account for this server.

NOTE: The Plug-in for VMware supports an ESXi Server that is in Lockdown mode if the host is managed by a vCenter Server.

To add a server:

- 1 In the Navigation pane, click **Create Backup Job**.
— or —
In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Create Backup Jobs**.
- 2 Click **Create New** next to the **Selections** list to open the **NetVault Backup Selections** page.
- 3 Open the NetVault Backup Client on which the plug-in is installed, and select **VMware Plugin**.
- 4 In the **Actions** list, click **Add Server**.
- 5 Configure the following settings:

Table 1. Add Server

Option	Description
Server Name	<p>In a standalone server setup, type a name for the ESXi Server.</p> <p>In a cluster setup, type a name for the vCenter Server.</p> <p>The server name must be unique. If you do not provide a name, the plug-in uses the Server Address as the server name.</p> <p>You cannot change the server name once it is configured.</p> <p>NOTE: When adding an ESXi Server or a vCenter Server to the plug-in installed on different clients, specify the same server name on all clients. If you use different names, the plug-in reports an error, “Server not found,” when you try to restore a virtual machine backed up from one server to the same server with a different name. This issue occurs because the restore path is different.</p>
Server Address	<p>In a standalone server setup, type the Fully Qualified Domain Name (FQDN) of the ESXi Server.</p> <p>In a cluster setup, type the FQDN of the vCenter Server.</p> <p>You can also specify the IP address of the server, but using FQDN is the preferred method.</p>
User	<p>Specify a user account that can be used to log on to the configured server. The user account must have privileges to register or create virtual machines.</p>
Password	<p>Type the password for the user account.</p>

- 6 To close the dialog box, click **OK**.

The plug-in tries to log on to the server, and if successful, adds the server to the backup selection tree.

NOTE: If a non-default port is configured on the vCenter Server to listen for connections from the vSphere Client, the plug-in cannot communicate with the server.

Reconfiguring a server

- 1 Start the Backup Job Wizard, and click **Create New** next to the **Selections** list.
- 2 Open the NetVault Backup Server, and then open **VMware Plugin**.
- 3 Select the ESXi or vCenter Server.
- 4 In the **Actions** list, click **Edit Server**.
- 5 Reconfigure the applicable settings.

For more information about these settings, see [Add Server](#). In the **Edit Server** dialog box, the server name is displayed as a read-only setting.

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

Removing a server

- 1 Start the Backup Job Wizard, and click **Create New** next to the **Selections** list.
- 2 Open the NetVault Backup Server, and then open **VMware Plugin**.
- 3 Select the ESXi or vCenter Server.
- 4 In the **Actions** list, click **Remove Server**.
- 5 In the confirmation dialog box, click **OK**.

Supported transport modes

Plug-in for VMware supports the following methods for accessing virtual machine disks:

- [SAN mode](#)
- [HotAdd mode](#)
- [LAN mode](#)

SAN mode

To use the SAN transport mode, the plug-in must be installed on a physical machine.

The SAN transport mode supports virtual machine disks stored on Fibre Channel SAN, iSCSI SAN, or Serial Attached SCSI (SAS) storage devices. This mode allows you to offload backups to a physical NetVault Backup Client.

SAN mode requirements

- With SAN storage, the NetVault Backup Client must have read and write access to the LUNs that contain the VMFS volumes (datastores) with virtual disks or the virtual compatibility RDMs. The NetVault Backup Client must be added to the same fabric zone to which the ESXi Server belongs. Moreover, the same host mode (connection type) must be set on the NetVault Backup Client and the ESXi Server. For example, if you have an IBM array with the host mode set to LNXCL, the same setting must be used for the NetVault Backup Client.
- To access VMFS datastores through iSCSI, you can use iSCSI initiators that are enabled on the ESXi hosts. The iSCSI initiators allow the host to connect to the iSCSI storage device through dedicated hardware or standard network adapters. For more information about configuring your systems, see the VMware iSCSI setup documentation and vendor-specific documentation.

NOTE: During restores, the SAN transport mode offers the best performance on thick disks. On thin disks, the NBD and NBDSSL transport modes are faster than the SAN mode.
During SAN restores, disable CBT on the virtual machine.

HotAdd mode

To use the HotAdd transport mode, the plug-in must be installed in a virtual machine.

The HotAdd method involves creating a linked clone of the target virtual machine and attaching the virtual disks to the backup proxy, which allows the disks to be read locally. However, these operations incur some overhead on the ESXi host, so the HotAdd mode is not as efficient as the SAN mode.

The HotAdd mode supports all types of storage devices. This mode does not require you to expose SAN LUNs to the NetVault Backup Client.

HotAdd mode requirements

- To use the SCSI HotAdd mode, the backup proxy virtual machine must reside on the same datacenter as the target virtual machines—that is, the virtual machines that you want to back up.
- The ESXi Server that hosts the backup proxy must have access to the datastores on which the backup proxy and target virtual machines reside.
- The virtual machines must use SCSI controller; the HotAdd mode does not support IDE and SATA controllers. The HotAdd mode only supports LSI SCSI controller; it does not support Paravirtual SCSI controllers.

LAN mode

To use the LAN mode, the plug-in can be installed either on a physical machine or in a virtual machine.

The LAN mode uses the Network Block Device (NBD) or NBDSSL protocol to access virtual disks. The ESXi Server host reads data from the storage device, and sends the read data across a network channel to the plug-in. The NBD transport mode performs unencrypted data transfers and can be used if the ESXi Server and Plug-in for VMware reside on a secure isolated network. NBD is faster than NBDSSL, and takes fewer resources on the ESXi Server and backup proxy. The NBDSSL transport mode uses SSL to encrypt all data passed over the TCP connection and can be used to protect sensitive data.

The LAN transport mode supports all types of storage devices. You can use the LAN transport mode if the ESXi Server uses local storage devices or NAS to store its virtual machine disks.

LAN mode requirements

To use secure communication channels (NBDSSL), enable SSL certificate verification in your virtual environment. This option is disabled by default.

Configuring default settings

- 1 In the navigation pane, click **Create Backup Job**, and then click **Create New** next to the **Selections** list.
- 2 Open the NetVault Backup Client on which the plug-in is installed, and select **VMware Plugin**.
- 3 In the **Actions** list, click **Configure**.

NOTE: You can also configure the default settings from the **Change Settings** page:

- 1 In the Navigation pane, click **Change Settings**.
- 2 If the plug-in is installed on the NetVault Backup Server, click **Server Settings**.
—or—
If the plug-in is installed on a NetVault Backup Client, click **Client Settings**; in the clients table, select the client, and click **Next**.
- 3 Under **Plugins**, click **Plugin Options**.

- 4 Under **Plug-in for VMware**, configure the following settings:

Table 2. Default settings for Plug-in for VMware

Option	Description
Transport Mode	<p>Select the transport mode for accessing virtual machine disks. The supported modes are:</p> <ul style="list-style-type: none"> • SAN • HotAdd • NBD or NBDSSL <p>For more information about transport modes, see Supported transport modes.</p> <p>To use the most suitable transport mode automatically, select Auto. On new plug-in installations, the Auto transport mode is selected by default.</p>
Fallback Transport Mode	<p>In the Fallback Transport Mode list, select the transportation mode that is used when the primary transportation mode fails. The available options are nbd, nbdssl, and none. If no alternative is available, select none.</p> <p>On new plug-in installations, the fallback transport mode is set to nbd by default.</p> <p>On upgrade installations, the fallback transport mode is set to nbdssl if the mode was set to san or hotadd mode before upgrade.</p>
Default Inventory View	<p>Plug-in for VMware provides two View Types to browse the VMware inventory objects on the NetVault Backup Selections page:</p> <ul style="list-style-type: none"> • Hosts and Clusters: The Hosts and Clusters view is the default inventory view for the Plug-in for VMware. <ul style="list-style-type: none"> The Hosts and Clusters view provides a hierarchical view of the hosts, clusters, and their child objects. In a cluster setup managed by a vCenter Server, the virtual machines are displayed under the Cluster node. In a standalone ESXi Server setup, the virtual machines are displayed under the individual hosts. The Hosts and Clusters view does not display the folders created on the vCenter Server. • VMs and Templates: The VMs and Templates view provides a flat view of all virtual machines and templates in the inventory; the virtual machines are grouped by Datacenters and folders. <p>NOTE: To switch between the two views, select the ESXi or vCenter Server on the NetVault Backup Selections page, and in the Actions list, click Toggle Inventory View. This option is only available when the server node is open.</p>
Progress Statistics Update Interval	<p>This setting determines the interval at which the plug-in updates progress statistics on the Monitor Job page. The default value for this option is 10 seconds. To change the default interval for progress updates, type or select a new value. The progress interval is specified in number of seconds.</p> <p>NOTE: Frequent progress updates can have a negative impact on the performance of the plug-in.</p>
Read Block Size	<p>This setting specifies the number of disk sectors to read or write per operation. The default value is 65536 sectors (one sector = 512 bytes; 65536 sectors = 32MiB). Setting a large block size for read and write operations can improve backup performance.</p>

Table 2. Default settings for Plug-in for VMware

Option	Description
Enable virtual machine locking	<p>If a virtual machine is migrated—for example, through Storage vMotion—while a backup or restore job is in progress, it can cause a job to fail. It can also create orphaned virtual disks on the datastores.</p> <p>You can use this option to enable locking of virtual machines against Storage vMotion during backup and restore operations. The virtual machines are locked before a job starts and unlocked after the job completes.</p> <p>This check box is selected by default.</p>
Attempts to acquire lock	<p>This option specifies the maximum number of attempts that are made to lock a virtual machine against Storage vMotion.</p> <p>The default value for this option is 10.</p>
Continue without lock	<p>By default, the backup of a virtual machine fails if it cannot be locked against Storage vMotion.</p> <p>When you select this check box, the plug-in continues to try to back up the virtual machine after the attempts to acquire a lock fail.</p>
Working Directory	<p>The Working Directory is used for the following purposes:</p> <ul style="list-style-type: none">• To mount virtual machine volumes during file-level browse and backup operations.• To store the data temporarily while restoring legacy backup savesets created with Plug-in for VMware 1.x. <p>The default path for the Working Directory is <code><NetVault Backup home>\tmp</code> on Windows® and <code><NetVault Backup home>/tmp</code> on Linux®. To change the location, specify the full path. If you specify a path that does not exist, the plug-in automatically creates it on the NetVault Backup machine.</p>

5 To save the settings, click **OK** or **Apply**.

Enabling or disabling CBT on virtual machines

This section includes the following topics:

- [About CBT](#)
- [Enabling CBT on individual virtual machines](#)
- [Disabling CBT on a virtual machine](#)
- [Manually disabling CBT on a virtual machine](#)

About CBT

The VMware Changed Block Tracking (CBT) feature allows virtual machines to track changed disk sectors. When CBT is enabled on a virtual machine, a “Change ID” is assigned to each disk when a snapshot is generated for backups. The Change ID identifies the state of a virtual disk at a specific point in time. Subsequent snapshots capture only the blocks that have changed since the last snapshot.

CBT offers the following advantages:

- It allows Incremental and Differential image-level backups of virtual machines.

- It often reduces the backup size of Full image-level backups because only the used sectors of a virtual disk are backed up.

CBT is only supported on virtual machines that use virtual hardware 7 or later. CBT is not supported on virtual machines that use Physical compatibility Raw Device Mapping (RDM) virtual disks, Virtual compatibility RDM (Independent Disks), or Virtual disks attached to a shared virtual SCSI bus.

Enabling CBT on individual virtual machines

By default, CBT is disabled on a virtual machine. If you want to perform Incremental or Differential Backups of virtual machines, you must enable CBT for the Full Backup that serves as the base backup for subsequent Incremental and Differential Backups.

The Plug-in for VMware provides the following methods to enable CBT on virtual machines:

- **Enable CBT on specific virtual machines:** To enable CBT on specific virtual machines, you can use the **Enable Change Block Tracking** method, which can be accessed from the **Actions** list on the **NetVault Backup Selections** page. This section describes the procedure for enabling CBT on individual virtual machines.
- **Enable CBT on all virtual machines included in a backup:** To enable tracking automatically on all virtual machines that are included in a backup job, you can set the backup option “**Enable Changed Block Tracking for VM**.” For more information, see [Performing image-level backups](#), [Enable Changed Block Tracking for VM](#).

To enable CBT on a virtual machine:

- 1 Start the Backup Job Wizard, and click **Create New** next to the **Selections** list.
- 2 Open the NetVault Backup Client on which plug-in is installed, and then open **VMware Plugin**.
- 3 Open the ESXi or vCenter Server and the other applicable container nodes—for example, Datacenter, Cluster, Resource Pool, and other nodes—to display the target virtual machine.
- 4 Select the virtual machine.
- 5 In the **Actions** list, click **Enable Change Block Tracking**.
This option is available to only those virtual machines on which CBT is disabled.
- 6 After the virtual machine is reconfigured, a message is displayed. To close the dialog box, click **OK**.

Important notes

- To create the first Full Backup of a virtual machine with Changed Block Tracking enabled, the target virtual machine must be powered off to enable CBT. For more information, see <http://kb.vmware.com/kb/1031873>.
The virtual machine can be restarted when the backup job starts transferring data. During subsequent Full, Incremental, or Differential Backups, the virtual machine can remain powered on.
- At the time of enabling CBT, no snapshots should exist on the target virtual machine. For more information, see <http://kb.vmware.com/kb/1033816>.
- The re-configuration of **Changed Block Tracking** takes effect only after the virtual machine goes through a stun and unstun cycle. This cycle includes power-on, resume after suspend, migrate, or snapshot create, delete, or revert operations.
- Once you enable CBT for a virtual machine, all subsequent Full, Incremental, and Differential Backups are performed as CBT-based backups for that virtual machine. You do not have to enable CBT for each job.

Disabling CBT on a virtual machine

Once you enable CBT on a virtual machine, all subsequent Full, Incremental, and Differential Backups are performed as CBT-based backups for that virtual machine. To stop using CBT, you must disable this feature.

- 1 Start the Backup Job Wizard, and click **Create New** next to the **Selections** list.
- 2 Open the NetVault Backup Client on which plug-in is installed, and then open **VMware Plugin**.
- 3 Open the ESXi or vCenter Server and the other applicable container nodes—for example, Datacenter, Cluster, Resource Pool, and other nodes—to display the target virtual machine.
- 4 Select the virtual machine.
- 5 In the **Actions** list, click **Disable Change Block Tracking**.
This option is available to only those virtual machines on which CBT is enabled.
- 6 After the virtual machine is reconfigured, a message is displayed. To close the dialog box, click **OK**.

Manually disabling CBT on a virtual machine

To manually disable CBT on a virtual machine, use the following procedure.

- 1 Power off the virtual machine.
- 2 In the virtual machine's configuration (“**.vmx**”) file, edit the following entry and set it to **False**:
`ctkEnabled = "False"`
- 3 In the “**.vmx**” file for each virtual disk, edit the following entry and set it to **False**:
`scsiox.ctkEnabled = "False"`
- 4 Create and delete a snapshot for the virtual machine to complete a stun and unstun cycle.
The “**.ctk**” files for the virtual machine are automatically deleted when you power on the virtual machine.

Enabling or disabling quiescing of virtual machines

This section includes the following topics:

- [About quiescing virtual machines](#)
- [Enabling quiescing of virtual machines](#)
- [Disabling quiescing of virtual machines](#)
- [Removing backup preferences for a virtual machine](#)

About quiescing virtual machines

A virtual machine can be quiesced to create consistent snapshots for backups. Quiescing is performed using the VMware VSS component, which is provided with VMware Tools. The VMware VSS component is automatically installed when you install VMware Tools on the virtual machines.

Quiescing is only supported on Windows®-based virtual machines. You can enable or disable quiescing of virtual machines.

Depending on the quiescing method used, the plug-in generates the following types of snapshots:

- **Crash-consistent snapshots:** Crash-consistent snapshots capture the volume state as it would be after a power outage. Such snapshots are not guaranteed to be free of incomplete I/O operations or data corruption. During a restore, there can be some data loss or corruption that can cause application failure.

① **NOTE:** On Linux®-based virtual machines, the plug-in always generates crash-consistent snapshots. You can install the VMware SYNC driver on Linux-based systems to freeze and thaw filesystem I/O and obtain a consistent snapshot that reduces the risk of data loss and data corruption after restore. However, creating quiesced snapshots using SYNC driver requires waiting for I/O to drain in the Guest OS. This activity can affect time-sensitive applications.

- **File-system-consistent snapshots:** For file-system-consistent snapshots, all file system I/O is temporarily frozen and dirty memory data is flushed to the disk before the snapshot is taken.

To create file-system-consistent snapshots, you must install the VMware VSS component on the virtual machines as part of VMware Tools.

- **VSS-based application-consistent snapshots:** For application-consistent snapshots, all application data in memory is written to the disk. Application-consistent snapshots reduce the amount of application crash recovery that you need to perform for VSS-enabled applications, like SQL Server®, Exchange, and SharePoint®, after restoring an image-level backup.

To create application-consistent snapshots, the following requirements must be met:

- You must install the VMware VSS component on the virtual machines as part of VMware Tools.
- The virtual machines should only use SCSI disks. Application-consistent snapshots are not supported on virtual machines that use IDE disks or dynamic disks.
- The virtual machines should have enough free SCSI slots as the number of disks.

Enabling quiescing of virtual machines

Quiescing is only supported on Windows®-based virtual machines. You can set the backup preferences to enable quiescing for all or specific virtual machines.

① **NOTE:** Quiescing is enabled by default.

- 1 Start the Backup Job Wizard, and click **Create New** next to the **Selections** list.
- 2 Open the NetVault Backup Client on which plug-in is installed, and then open **VMware Plugin**.
- 3 To enable quiescing for all virtual machines managed by an ESXi or vCenter Server, select the server, and in the **Actions** list, click **Set Backup Preferences**.

To enable quiescing for specific virtual machines, open the ESXi or vCenter Server and the other applicable container nodes—for example, Datacenter, Cluster, Resource Pool, and other nodes—to display the target virtual machines. Select the virtual machine, and in the **Actions** list, click **Set Backup Preferences**.

- 4 In the **Backup Preferences** dialog box, verify that the **Quiesce VM for snapshot** option is selected. This check box is selected by default. To quiesce a virtual machine before creating a snapshot, leave this check box selected. The setting for an individual virtual machine level takes precedence over the server-level setting.
- 5 To close the dialog box, click **OK**.

Disabling quiescing of virtual machines

On virtual machines with high CPU or IO load, the quiesce operation may time out and cause the backup job to fail. In such cases, you can disable quiescing for all or specific virtual machines. When you disable quiescing, the plug-in uses crash-consistent snapshots to perform backups.

- 1 Start the Backup Job Wizard, and click **Create New** next to the **Selections** list.
- 2 Open the NetVault Backup Client on which plug-in is installed, and then open **VMware Plugin**.
- 3 To disable quiescing for all virtual machines managed by an ESXi or vCenter Server, select the server, and in the **Actions** list, click **Set Backup Preferences**.

To disable quiescing for specific virtual machines, open the ESXi or vCenter Server and the other applicable container nodes—for example, Datacenter, Cluster, Resource Pool, and other nodes—to display the target virtual machines. Select the virtual machine, and in the **Actions** list, click **Set Backup Preferences**.

- 4 In the **Backup Preferences** dialog box, clear the **Quiesce VM for snapshot** option.
Clearing this check box creates a snapshot without quiescing a virtual machine. The setting for an individual virtual machine level takes precedence over the server-level setting.
- 5 To close the dialog box, click **OK**.

Removing backup preferences for a virtual machine

By removing backup preferences for a virtual machine, you can ensure that the plug-in automatically applies the server preferences to the individual virtual machine. If you enable or disable quiescing at the server level, the same setting is applied to the virtual machine for which the configuration has been removed.

- 1 Start the Backup Job Wizard, and click **Create New** next to the **Selections** list.
- 2 Open the NetVault Backup Client on which plug-in is installed, and then open **VMware Plugin**.
- 3 Open the ESXi or vCenter Server and the other applicable container nodes—for example, Datacenter, Cluster, Resource Pool, and other nodes—to display the target virtual machines.
- 4 Select the virtual machine.
- 5 In the **Actions** list, click **Remove Backup Preferences**.

Unlocking a virtual machine

You can use the following procedure to unlock a virtual machine from the WebUI.

- 1 Start the Backup Job Wizard, and click **Create New** next to the **Selections** list.
- 2 Open the NetVault Backup Client on which plug-in is installed, and then open **VMware Plugin**.
- 3 Open the ESXi or vCenter Server and the other applicable container nodes—for example, Datacenter, Cluster, Resource Pool, and other nodes—to display the target virtual machine.
- 4 Select the applicable virtual machine, and in the **Actions** list, click **Unlock Virtual Machine**.

Defining a backup strategy

- [About backup methods and types](#)
- [Data backed up and restored for different disk types](#)
- [Backup and recovery strategy](#)

About backup methods and types

The Plug-in for VMware supports image-level and file-level backup methods:

- [Image-level backups](#)
- [File-level backups](#)

Image-level backups

Image-level backups use the VMware snapshot technology to provide point-in-time image of virtual machines. You can use these backups to perform the following types of recoveries:

- Recover a full virtual machine to a previous known state.
- Restore one or more virtual disks for a virtual machine.
- Restore individual files and directories to a specified location.
- Restore virtual machine disk and configuration files to a specified location.

Image-level backups are supported on Linux[®]- and Windows[®]-based virtual machines. These backups can be performed with or without CBT.

Image-level backups with CBT

When CBT is enabled on a virtual machine, the following backup types are supported:

- **Full Backup:** A Full Backup provides a backup of all allocated sectors on a virtual disk. Full Backups take longer to complete and consume more backup media. These backups serve as the base backup for future Incremental and Differential image-level backups.
- **Differential Backup:** A Differential Backup provides a backup of disk sectors that have changed since the last Full Backup. Differential Backups speed up recovery since the plug-in is only required to restore two savesets.
- **Incremental Backup:** An Incremental Backup provides a backup of disk sectors that have changed since the last Full, Differential, 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.

Image-level backups without CBT

Only Full image-level backups are supported on virtual machines that do not use CBT.

File-level backups

The file-level backup method is only available to Windows®-based virtual machines. File-level backups allow you to restore a single file or a set of files. You can use these backups to recover from data loss due to user errors, data corruption, or accidental deletion of files.

NOTE: Only the Windows version of the Plug-in for VMware supports file-level backups of virtual machines.

The Plug-in for VMware uses NetVault Backup Plug-in for FileSystem (Plug-in for FileSystem) for file-level backups.

File-level backup types

The Plug-in for VMware supports the following types of file-level backups:

- **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
- **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.
- **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.

Data backed up and restored for different disk types

The Plug-in for VMware includes the following data in CBT-enabled backups and restores.

Table 3. Data backed up and restored for different disk types

Disk Type	CBT-enabled Full Image Backups	Incremental/Differential Image Backups	Restore of CBT-enabled Full Image Backups	Restore of Incremental/Differential Image Backups
Thin (Required space supplied and zeroed out on demand.)	Only used disk sectors are backed up.	Only changed disk sectors are backed up.	Only used sectors are restored.	Only used disk sectors are restored. Each sector is restored only once.
Zeroed Thick (All space allocated at creation time, unused part zeroed out on first write.)	Only used disk sectors are backed up.	Only changed disk sectors are backed up.	Only used sectors are restored.	Only used disk sectors are restored. Each sector is restored only once.
Eager Zeroed Thick (All space allocated and zeroed out at creation time.)	All disk sectors are backed up.	Only changed disk sectors are backed up.	The entire disk is restored.	The entire disk is restored. Each sector is restored only once.
Virtual Compatibility RDM (Raw Device Mapping)	All disk sectors are backed up.	Only changed disk sectors are backed up.	The entire disk is restored.	The entire disk is restored. Each sector is restored only once.
Network File System (NFS)	All disk sectors are backed up.	Only changed disk sectors are backed up.	The entire disk is restored.	The entire disk is restored. Each sector is restored only once.

- ① **NOTE:** Regardless of the disk provisioning type, image-level backups without CBT always back up the full disk size. Thus, backup of a 20GB thin-provisioned disk that only uses 2GB results in a 20GB backup. When you restore a non-CBT backup, all blocks are restored and allocated. After a disk is restored from a non-CBT backup, subsequent CBT-enabled Full Backups back up the entire disk.

RDM disks

The following table describes how the plug-in handles RDM (Raw Device Mapping) disks.

Table 4. Considerations for RDM disks

RDM compatibility mode	VADP-based backups and restores
Physical compatibility mode	The disk is not backed up. Warning messages are written to the NetVault Backup binary logs and job logs.
Virtual compatibility mode (Independent disks)	The disk is not backed up. Warning messages are written to the NetVault Backup binary logs and job logs.
Virtual compatibility mode	<p>The disk is backed up, but the data can only be restored to a flat file. Warning messages about the change in type are written to the NetVault Backup binary logs and job logs.</p> <p>Optionally, you can omit the RDM disks in virtual compatibility mode during a restore.</p> <p>To omit the RDM disks in virtual compatibility mode:</p> <ol style="list-style-type: none">1 Open the <code>vmware.cfg</code> file in a text editor. You can find this file in <code><NetVault Backup home>\config</code> on Windows® and <code><NetVault Backup home>/config</code> on Linux®.2 Edit the following entry and set the value to <code>false</code>. <pre>[Custom:RestoreRDMDisks] Value=False</pre>3 Save the file.

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.

Using the image-level backup method

- [Performing image-level backups](#)
- [Restarting a backup job](#)
- [Resetting CBT on a virtual machine](#)
- [Monitoring job progress](#)

Performing image-level backups

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

— or —

In the **Navigation** pane, click **Guided Configuration**, and then 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 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 To open the **NetVault Backup Selections** page, click **Create New**.
- b Open the NetVault Backup Client on which the plug-in is installed, and then open **VMware Plugin**.
- c Open the applicable VMware® ESXi™ or VMware® vCenter™ Server.

Depending on your VMware setup and inventory view type, open the applicable container nodes.

- ① **NOTE:** To switch between the two views (**Hosts and Clusters** and **VMs and Templates**), select the ESXi or vCenter Server, and in the **Actions** list, click **Toggle Inventory View**. This option is only available when server node is open.

- d Select the data that you want to back up:
 - **Back up all virtual machines within a container:** Select the container node. Data selection is possible at all levels, starting from the Datacenter node. For example, to back up all virtual machines hosted on an ESXi server, select the host node; in the **VMs and Templates** view, select the corresponding folder node.
 - **Back up individual virtual machines:** Open the applicable container nodes, for example, Datacenter, Cluster, and Resource Pool, and select the virtual machines that you want to back up. In the **VMs and Templates** view, open the folder node, and select the virtual machines that you want to back up.

You can also select the container node, and then clear the check marks for the virtual machines that you do not want to back up.

- **Back up individual virtual disks:** Open the virtual machine node, and select the disks that you want to back up. The virtual disks are named “Hard Disk 1,” “Hard Disk 2,” ... “Hard Disk n.” The plug-in only lists those disks for which it can generate a snapshot.

You can also select the container node—for example, Datacenter, Resource Pool, ESXi Server, or folder—or the individual virtual machines, and specify the disk type—system or data—in the Backup Options Set. For more information about this setting, see [Virtual machine disk selection options](#).

The virtual machine configuration files, “.nvram”, and log files are always backed up whether you select the entire virtual machine or its individual disks.

- e Click **Save**, and in the **Create New Set** dialog box, type a name for the set.

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.

- f To close the dialog box, click **Save**.

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

① **IMPORTANT:** On Windows, when you create a Backup Options Set for image-level backups, you must use the default set “**Default Backup Options – VMware Plugin – Windows**” as the template. If you use any other set as the template, the backup may fail.

- a To open the **VMware Plugin Backup Options** page, click **Create New**.
- b Under **Backup Type**, select one of the following options.

Table 5. Image-level backup types for virtual machines

Option	Description
Full	Select this option to back up all allocated sectors on a virtual disk.
Incremental	Select this option to back up disk sectors that have changed since the last Full, Differential, or Incremental Backup.
Differential	Select this option to back up disk sectors that have changed since the last Full Backup.

For more information about image-level backup types, see [Image-level backups](#).

① **IMPORTANT:** After reverting to a snapshot that is older than its last Incremental Backup, you must perform a Full Backup of that virtual machine before performing Incremental Backups again. If you perform an Incremental Backup, the job reports a filefault error and fails. For more information, see <http://kb.vmware.com/kb/1021607>.

- c Under **VM disk selection options**, select one of the following options.

The disk selection option is applied to all virtual machines that you have selected explicitly by selecting the corresponding nodes or implicitly by selecting the container node—for example, ESXi Server, Resource Pool, Datacenter, or folder. This option is not applied to virtual machines for which you have explicitly selected one or more virtual disks in the selections tree.

Table 6. Virtual machine disk selection options

Option	Description
Back up all disks	To back up all available virtual disks for the selected virtual machines, use this option. NOTE: The plug-in only backs up those disks for which it can generate a snapshot.
Exclude boot disk	To back up only the data disks and exclude the boot disks for the selected virtual machines, use this option.
Exclude data disks	To back up only the boot disk and exclude the data disks for the selected virtual machines, use this option. NOTE: To identify the boot disks, the plug-in only considers the active partitions on the MBR disk. Moreover, the plug-in only considers the first disk that meets the active boot partition criteria as the boot disk. The plug-in does not handle machines with multiple operating systems.

d Under **Other Options**, configure the following settings:

Table 7. Other backup options for Plug-in for VMware

Option	Description
Enable Changed Block Tracking for VM	<p>By default, CBT is disabled for a virtual machine. If you want to perform Incremental or Differential Backups of virtual machines, you must enable CBT for the Full Backup that serves as the base backup for subsequent Incremental and Differential Backups.</p> <p>To enable CBT on all virtual machines included in the backup job, select this check box. (You can also enable CBT on specific virtual machines by using the Enable Change Block Tracking method. For more information, see Enabling CBT on individual virtual machines.)</p> <p>If you select this check box and the plug-in fails to change this setting on a virtual machine, a warning message is logged.</p> <p>If you do not select this check box, the plug-in does not change the CBT setting on the virtual machines during backups. Depending on whether CBT is enabled or disabled, the plug-in uses the appropriate backup method (CBT-based Full, Incremental, or Differential or non-CBT Full Backup) to back up the virtual machines.</p> <p>Note the following:</p> <ul style="list-style-type: none"> To create the first Full Backup of a virtual machine with Changed Block Tracking enabled, the target virtual machine must be powered off to enable CBT. For more information, see http://kb.vmware.com/kb/1031873. The virtual machine can be restarted when the backup job starts transferring data. During subsequent Full, Incremental, or Differential Backups, the virtual machine can remain powered on. At the time of enabling CBT, no snapshots should exist on the target virtual machine. For more information, see http://kb.vmware.com/kb/1033816. Once you enable CBT for a virtual machine, all subsequent Full, Incremental, and Differential Backups are performed as CBT-based backups for that virtual machine. You do not have to select this option for each job. <p>To stop using CBT, you must disable this feature. For more information, see Disabling CBT on a virtual machine or Manually disabling CBT on a virtual machine.</p>

Table 7. Other backup options for Plug-in for VMware

Option	Description
Perform File Level Indexing	<p>When you perform an image-level backup, the plug-in automatically generates file-level indexes for all supported volumes that are included in the backup snapshot. File-level indexing allows you to restore individual files and directories from Full, Incremental, and Differential image-level backups of virtual machines.</p> <p>File-level indexing is available to volumes that use the following file systems:</p> <ul style="list-style-type: none"> • Windows®: NTFS • Linux® and UNIX®: EXT2, EXT3, and EXT4 <p>The Plug-in for VMware 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.</p> <p>NOTE: The current version of the plug-in does not support Windows Server® 2012 ReFS (Resilient File System) and Striped Disks.</p> <p>File-level indexing does not affect the backup size. However, it increases the backup index size and the total backup time. The amount of time taken to perform file-level indexing depends on several factors, including the number of files, fragmentation of files on the volumes, network traffic, and load on the ESXi or vCenter Server.</p> <p>If you do not want to use image-level backups for file-level restores, clear this check box.</p>
Enable Active Block Mapping	<p>The Active Block Mapping (ABM) technology provides filters to remove unused blocks during backups. Removing the unused blocks reduces the backup size and the amount of data transferred over the network. You can use ABM with CBT to back up only active and changed blocks during Incremental and Differential Backups.</p> <p>ABM is supported on the NTFS file system on basic disks and EXT file systems.</p> <p>ABM is disabled by default. Selecting this check box enables ABM, which allows the plug-in to back up only those sections of the virtual disks that are active. ABM scans the disk and detects inactive blocks; these blocks are then omitted during backups.</p> <p>NOTE: ABM also removes deleted data. When you restore a virtual machine that was backed up using ABM, you cannot perform undelete operations for that virtual machine. Without Active Block Mapping, CBT includes the deleted blocks.</p> <p>If a backup job includes any unsupported disk types, this option is ignored for those disks.</p>
Remove snapshots from previous backups	<p>When you run a backup job, the plug-in creates a snapshot named “BKB_SNAP” on the virtual machine to back up the selected data. Regardless of whether the backup completes successfully or fails, the plug-in removes the snapshot on job completion. The plug-in also removes the snapshot if the job is aborted or the parent or child processes are terminated for any reason. However, if a job exits abnormally, the cleanup process may fail to remove the snapshot. To address such cases, the plug-in provides an option to delete the snapshot when you run the job next time.</p> <p>You can select the Remove snapshots from previous backups check box to remove any existing snapshots when you run the current job. Only the snapshots named “BKB_SNAP” are removed from the virtual machines. This option does not remove any other snapshots that exist on the virtual machine.</p>

Table 7. Other backup options for Plug-in for VMware

Option	Description
Maximum Number of Parallel Streams	<p>By default, the plug-in generates a single data stream for a backup job, and backs up the selected virtual machines in a sequential manner. To increase throughput and reduce overall backup time for image-level backups, you can configure the plug-in to perform parallel backups of multiple virtual machines.</p> <p>This setting determines the maximum number of parallel data streams that can be generated for an image-level backup job. For example, if 10 virtual machines are included in a job, and you set this parameter to four, the plug-in tries to back up four virtual machines in parallel.</p> <p>The actual number of parallel streams for a job depends on the following factors:</p> <ul style="list-style-type: none"> • The number of backup devices or streams available to the backup job. For example, if you set this parameter to 4 and only two tape drives are available or a NetVault SmartDisk can only support two concurrent streams, the plug-in only processes two virtual machines concurrently. • The number of virtual machines included in the job. For example, if you set this parameter to 4, and the number of virtual machines selected for backup is 3, the plug-in creates only three data streams. <p>For parallel backups, the plug-in generates a parent process that coordinates the overall backup and individual child processes that perform the actual task of backing up the virtual machines. The maximum child processes that can be generated for a backup job is equal to the Maximum Number of Parallel Streams configured for the job. The parent and the child processes are all created on the NetVault Backup Client on which the plug-in is running.</p> <p>When a child process acquires a backup device and creates a backup stream, it is assigned a task to back up a virtual machine. After the task completes, the process is assigned the next task if there are any more virtual machines to be backed up. Each task is assigned a Task ID. A snapshot is generated only when a virtual machine is assigned to a child process.</p> <p>When configuring Maximum Number of Parallel Streams, consider the following:</p> <ul style="list-style-type: none"> • Verify that the value configured for this option does not exceed the number of backup devices or streams accessible to a job. If the available devices are insufficient, more than one child process may try to access and write to the same device, causing significant increase in the overall backup time. • The load of running multiple processes on a NetVault Backup Client can have a negative impact on performance. • The load of backing up multiple virtual machines from the same datastore can increase the I/O activity on the datastore. • The load of backing up multiple virtual machines that are managed by the same ESXi Server can increase the load on the host if you are using the LAN Transport mode (nbd/nbdssl) for data transfer. <p>NOTE: Only a single child process is used to back up the multiple disks of a virtual machine even if more data streams are available. The child process creates a snapshot of the whole virtual machine and retains the snapshot until the job completes regardless of whether the whole virtual machine or its individual disks are backed up.</p>

Table 7. Other backup options for Plug-in for VMware

Option	Description
Automatic diagnose virtual machine on backup failure	<p>NOTE: This setting is not applied to backup jobs created with plug-in versions earlier than 2.5. For such jobs, the plug-in generates a single stream and backs up the virtual machines sequentially.</p> <p>To run predefined tests that can identify the cause of a virtual machine backup failure, select this check box.</p> <p>For more information about the Diagnose method, see Diagnosing virtual machine issues.</p> <p>You can access the Diagnostic Results dialog box from the View Logs page. The details are stored as log context objects.</p>
Enable Restartable backups	<p>This option allows you to restart a job that completes with some failed virtual machines. The restarted instance backs up only those virtual machines that failed previously; virtual machines that are backed up successfully are not included in the restarted instance.</p> <p>When a restartable backup completes with some failed virtual machines, the plug-in generates a backup index for the completed virtual machines and sets the job status to Job Stopped. The log messages and log context show which virtual machines have failed for the job. When you restart the job later, the plug-in runs an Incremental Backup job to back up the failed virtual machines.</p> <p>You can restart a job from the Job Status page. For more information about this method, see Restarting a backup job. A restarted instance does not back up virtual machines that are added to the host after the initial job is run.</p> <p>NOTE: If all virtual machines fail for a job, the job status is set to Backup Failed. You cannot restart a failed backup job.</p>

- e Click **Save**, and in the **Create New Set** dialog box, type a name for the set.

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.

- f To close the dialog box, click **Save**.

- 5 Select or create the Schedule Set, Target Set, and Advanced Options Set.

For more information about these sets, see the *Dell NetVault Backup 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.

- 6 To submit the job for scheduling, click **Save & Submit**.

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.

You can monitor the job progress from the **Job Status** page and view the logs from the **View Logs** page.

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

Additional notes

- When you select a container node (for example, Datacenter, Resource Pool, or ESXi Server), any new virtual machine that is added to the host after a job has been defined is automatically included in the backup. Similarly, when you remove a virtual machine from the host, it is automatically excluded from the backup. You do not have to modify the job definition each time a virtual machine is added to or removed from the host.
- If a selected virtual machine is moved to a different host using Storage vMotion after a backup job has been defined, it is backed up as part of that job only if the following requirements are met:
 - A vCenter Server has been added to the plug-in, and the access to the virtual machine is being authenticated through the vCenter Server.
 - The ESXi Server Host to which the selected virtual machine has been moved to is also controlled by the same vCenter Server.

If these requirements are not met, the plug-in fails to back up the virtual machine that has been moved.

- When backing up virtual machines that are protected using VMware vSphere[®] Fault Tolerance (vSphere FT), note the following:
 - The vSphere FT group must be managed under a vCenter Server by VMware[®] ESXi[™] hosts that use version 6.0 build number 4192238 or later.
 - You can only select the primary virtual machines in an FT group for backups. The plug-in does not let you select the secondary virtual machines.
 - The plug-in displays the primary virtual machine in an FT group for selection only when they are protected using vSphere FT. If the virtual machine is protected using legacy fault tolerance, you cannot select it.
 - For container-level selections, only primary virtual machines in an FT group are backed up
 - Secondary virtual machines are excluded from the backup. If one is encountered, the plug-in generates the following log message:

```
Virtual machine '<VM_name>' with secondary role
in its fault tolerance group will not be backed up.
```
 - Primary virtual machines using Legacy FT are also not backed up. If one is encountered, the plug-in generates the following log message and completes with a Warning: Cannot backup virtual machine while legacy fault tolerance is enabled.
 - If a backup job includes an FT machine and a failover occurs, the job automatically protects the new primary virtual machine the next time the job runs.
 - If a failover occurs while a backup job is running, the backup job might fail because the backup snapshot cannot be created. VADP prevents the creation of snapshots during failover of a virtual machine with fault tolerance enabled. The log displays the following message:

```
Task to create VM snapshot failed with reason
'The operation is not allowed in the current state.'
```

If this issue occurs, run the backup job again after the failover process is finished.
 - Because VMware does not support CBT for virtual machines with FT enabled, all sectors on the disks are always backed up for such a virtual machine even if **Enable CBT...** is selected. This behavior is true for both Full and Incremental Backup jobs.
 - When you restore a virtual machine that had FT enabled, the plug-in does not configure FT for the machine after recovery. You must enable FT on the restored machine after recovery is completed. The plug-in generates the following log message and completes the restore job with a Warning: Fault tolerance was configured for the virtual machine at backup time; to continue using the feature, please re-configure it after recovery.

- When you restore data for selective virtual disks of an existing virtual machine, the plug-in fails the restore if the target virtual machine is configured for FT. The plug-in generates the following log message: Disk-level restores are currently not supported for fault tolerant virtual machines.
- VMware does not support file-level mounting for virtual machines with FT enabled.
- Although the plug-in does not support file-level (.vmdk) restores of virtual machines that have FT enabled, you can restore files at the guest filesystem level.
- You can select virtual machines that are part of a Virtual Application (vApp) along with the other virtual machines. Currently, the selection tree does not provide any indicators to distinguish between such virtual machines from other virtual machines that are not part of a vApp.

If you select a vApp node, only the virtual machines contained in the vApp are backed up. The metadata for the vApp is not included in the backup.

- If you add a new virtual machine to the Backup Selection Set—or to an ESXi Server that is included in 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 virtual machine if you run Incremental or Differential Backup.
 - Creates a non-CBT Full Backup, unless CBT is already enabled for the virtual machine or you have selected the **Enable Change Block Tracking for VM** check box for the Incremental or Differential Backup.
 - Runs subsequent Incremental Backups of the new virtual machine as Incremental if CBT is enabled and a CBT-based Full Backup has been performed for the virtual machine.
 - Runs subsequent Differential Backups of the new virtual machine as Full Backups because these backups are based on the last Full Backup created for that Backup Selection Set.

Example:

- 1 Select any one virtual machine, for example, MyVM1, and create a Backup Selection Set, for example, MySelectionSet.
- 2 Perform a CBT-enabled Full Backup using MySelectionSet.
- 3 Modify MySelectionSet, and add a virtual machine, for example, MyVM2.
- 4 Perform CBT-enabled Incremental and Differential Backup using MySelectionSet.
- 5 Perform another CBT-enabled Incremental and Differential Backup using MySelectionSet.

Result:

- When you select Incremental Backup, the plug-in creates a Full Backup of MyVM2 in [Step 4](#) and an Incremental Backup of that virtual machine in [Step 5](#).
- When you select Differential Backup, the plug-in creates a Full Backup of MyVM2 in [Step 4](#) and [Step 5](#) because these backups are based on the last Full Backup created in [Step 2](#).
- When using CBT for backups, virtual machine migration using Storage vMotion or VMware vSphere® Storage DRS™ is not recommended. For more information, see <http://kb.vmware.com/kb/2048201>.
- If you perform a Full Backup of a Thick Provisioned Lazy Zeroed disk without enabling CBT, the empty sectors are converted to actual zeros during backup. When you restore the backup, the disk type is changed to Eager Zeroed.
- After backing up the data, the plug-in writes the backup index to the NetVault Database. If an error occurs during this step, for example, if there is a file transfer error or the server runs out of disk space, the plug-in reports the following error:

Log message: Failed to write index of backup to the database.

Log context: The index for this backup may be retrieved and added to the database by scanning the backup media.



In such cases, you can scan the backup media to recover the index. However, before importing the index, you must review the log messages to verify that no other errors were reported during the backup.

Backup selection tree icons

Table 8. Backup selection tree icons

Icon	Description
	vCenter Server
	Datacenter Server
	ESXi Server Cluster
	Closed folder
	Open folder
	ESXi Server
	ESXi Server in maintenance mode
	Inaccessible ESXi Server
	Virtual Appliance (vApp)
	Resource pool
	Virtual machine (powered on)
	Inaccessible virtual machine
	Mounted virtual machine
	Suspended virtual machine
	Powered off virtual machine
	Virtual machine (powered on, CBT enabled)
	Inaccessible virtual machine (CBT enabled)
	Mounted virtual machine (CBT enabled)
	Suspended virtual machine (CBT enabled)
	Powered off virtual machine (CBT enabled)
	Primary virtual machine (powered on) in a Fault Tolerant group
	Inaccessible primary virtual machine
	Suspended primary virtual machine
	Powered off primary virtual machine
	Secondary virtual machine in a Fault Tolerant group
	Inaccessible secondary virtual machine

Table 8. Backup selection tree icons

Icon	Description
	Suspended secondary virtual machine
	Powered off secondary virtual machine

Restarting a backup job

The Plug-in for VMware provides the ability to restart a job to back up only those virtual machines that failed previously. To use this method, you must configure the backup option **Enable Restartable backups** for the job. For more information about this option, see [Performing image-level backups](#). When a restartable backup completes with some failed virtual machines, the plug-in generates a backup index for the completed virtual machines and sets the job status to **Job Stopped**. You can restart this job from the **Job Status** page.


 | **NOTE:** The **Restart** method does 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**.

When you restart a job, the plug-in runs an Incremental Backup job to back up the failed virtual machines. The restarted instance does not back up virtual machines that are added to the host after the initial job is run.

All restarted instances appear as a single saveset on the **Create Restore Job – Choose Saveset** page.

 | **IMPORTANT:** You can restart the job more than once; however, if you try to restart an instance after running the subsequent Full or Incremental Backup for that backup sequence, the job reports an index conflict error.

Resetting CBT on a virtual machine

During Incremental or Differential backups, if the plug-in is unable to determine the changed disk sectors for a virtual machine, it reports an error—“Failed to determine disk areas to back up”—and backs up the entire disk. When this error occurs, use the following procedure to reset change block tracking for the virtual machine. Before proceeding, delete any existing snapshots of the target virtual machine.

- 1 Start the Backup Job Wizard, and click **Create New** next to the **Selections** list.
- 2 Open the NetVault Backup Client on which plug-in is installed, and then open **VMware Plugin**.
- 3 Open the ESXi or vCenter Server and the other applicable container nodes—for example, Datacenter, Cluster, Resource Pool, and other nodes—to display the target virtual machine.
- 4 Select the target virtual machine, and in the **Actions** list, click **Reset Change Block Tracking**.
This option is available to only those virtual machines on which CBT is enabled.
- 5 After the virtual machine is reconfigured and a message is displayed, click **OK** to close the dialog box.

Important notes

- When CBT is being reset, no snapshots should exist for the virtual machine. Otherwise, the reset operation fails.

- If CBT is reset for a virtual machine that is powered off, Dell recommends that you power on the virtual machine before running a CBT-enabled backup for that virtual machine. If you perform a CBT-enabled backup in the powered off state, the backup may report the following error and fail:

Log message: Failed to determine disk areas to back up.

Log context: A specified parameter was not correct. deviceKey.

- After CBT is reset for a virtual machine, the next Incremental Backup backs up all blocks for that virtual machine. Subsequent Incremental Backups only back up the disk sectors that have changed.
- If an error message is displayed while CBT is being reset, check the “Recent Tasks” window on the vSphere Client for any messages that explain why the request failed.

Monitoring job progress

- 1 In the Navigation pane, click **Job Status**.
- 2 In the list of jobs, select the applicable job, and click **Monitor**.
- 3 On the **Monitor Job** page, you can view the following information.
 - **Job details:** This area displays the Job ID, title, phase, instance, client, plug-in, start time, expected completion time, run count, duration, size, and status.
 - **Data transfer chart:** This area displays the data transfer chart.
 - **Job logs:** This area displays the log messages.

Using the file-level backup method

- Performing file-level backups
- Removing the snapshot and mount folder manually

Performing file-level backups

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

— or —

In the **Navigation** pane, click **Guided Configuration**, and then 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 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 To open the **NetVault Backup Selections** page, click **Create New**.
- b Open the NetVault Backup Client on which the plug-in is installed, and then open **VMware Plugin**.
- c Open the applicable VMware® ESXi™ or VMware® vCenter™ Server.

Depending on your inventory view type, do the following:

- **Hosts and Clusters Inventory View:** Open the Datacenter, Cluster, Resource Pool, and other nodes to display the available virtual machines.
- **VMs and Templates View:** Open the Datacenter and folder nodes to display the available virtual machines.

- d Select the target virtual machine, and in the **Actions** list, click **Mount**.

The plug-in tries to obtain a snapshot of the virtual disk files and mount it on the NetVault Backup Client. Depending on the file size of the virtual machine disk files, the mount operation can take a few minutes to complete. When the snapshot is mounted successfully, the following changes occur:

- A folder is created in the **Working Directory** (set in the **Configure** dialog box). This folder is assigned the same name as the selected virtual machine.
- A **Drives** node is added to the selections tree. This node is displayed under selected virtual machine.

① **NOTE:** If the boot disk is not the first disk among the disks attached to a virtual machine when they are arranged in the **controller:device** order (that is, **ide0:0**, **ide0:1**, **scsi0:0**, **scsi0:1**, and so on), the mounting operation might fail for the virtual machine.

- e Open the **Drives** node to display the available drives.

- f Select the drives, or open the nodes further to display the directory tree, and select the files and directories that you want to back up.
 - ① **IMPORTANT:** After mounting a virtual machine, if you select the virtual machine node, the job fails. To include all drives in a job, you must select each drive individually.
 - g Click **Save**, and in the **Create New Set** dialog box, type a name for the set.
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.
 - h To close the dialog box, click **Save**.
- 4 In the **Plugin Options** list, select an existing Backup Options Set, or complete the following steps to create a set:
- ① **IMPORTANT:** On Windows, when you create a Backup Options Set for file-level backups, you must use the default set “**Default Backup Options for 'File System' – VMware Plugin - Windows**” as the template. If you use any other set as the template, the backup may fail.
 - a To open the **File System Plugin Backup Options** page, click **Create New**.
 - b Configure the following settings:

Table 9. Backup options for Plug-in for FileSystem

Option	Description
Backup Type	Select the applicable option from the following: <ul style="list-style-type: none"> • Full • Incremental • Differential For more information about file-level backup types, see File-level backups .
Create Dump Type Backup	To create Dump Type Incremental or Differential Backups, select this check box. With these backups you can only restore those files that were backed up in a particular Incremental or Differential Backup.
Check For Files Being Modified During Backup	To mark the files that are being modified during a backup as “in flux,” 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.
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. 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 the <i>Dell NetVault Backup Plug-in for FileSystem User’s Guide</i> .
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, and includes details such as file size, date modified, and file type.

- c Click **Save**, and in the **Create New Set** dialog box, type a name for the set.

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.

- d To close the dialog box, click **Save**.

- 5 Select or create the Schedule Set, Target Set, and Advanced Options Set.

For more information about these sets, see the *Dell NetVault Backup Administrator's Guide*.

- 6 To submit the job for scheduling, click **Save & Submit**.

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.

You can monitor the job progress from the **Job Status** page and view the logs from the **View Logs** page.

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

Backup selection tree icons

Table 10. Backup selections tree icons

















Icon	Description
	vCenter Server
	Datacenter Server
	ESXi Server Cluster
	Closed folder
	Open folder
	ESXi Server
	ESXi Server in maintenance mode
	Inaccessible ESXi Server
	Virtual Appliance (vApp)
	Resource pool
	Virtual machine (powered on)
	Inaccessible virtual machine
	Mounted virtual machine
	Suspended virtual machine
	Powered off virtual machine
	Virtual machine (powered on, CBT enabled)

Table 10. Backup selections tree icons

Icon	Description
	Inaccessible virtual machine (CBT enabled)
	Mounted virtual machine (CBT enabled)
	Suspended virtual machine (CBT enabled)
	Powered off virtual machine (CBT enabled)
	Primary virtual machine (powered on) in a Fault Tolerant group
	Inaccessible primary virtual machine
	Suspended primary virtual machine
	Powered off primary virtual machine
	Secondary virtual machine in a Fault Tolerant group
	Inaccessible secondary virtual machine
	Suspended secondary virtual machine
	Powered off secondary virtual machine

Removing the snapshot and mount folder manually

When you mount a virtual machine for file-level backup or browse operation, the following events occur:

- The plug-in creates a snapshot named “**BKB_SNAP**” on the virtual machine.
- The plug-in creates a mount folder for the virtual machine in the Working Directory. This folder is assigned the same name as the virtual machine.

When you unmount the virtual machine, the cleanup process automatically removes the mount folder and snapshot. In a normal scenario, do not remove them manually.

If the plug-in fails to remove the mount folder or snapshot for any reason, a subsequent mount operation for the same virtual machine fails and the error message “A stale mount was found” is displayed. For example, if the plug-in exits unexpectedly after mounting a virtual machine, the snapshot and mount folder are not removed. In this scenario, you must manually remove them. You must also complete these steps if a snapshot is deleted manually while a virtual machine is still mounted.

To remove the snapshot and mount folder manually:

- 1 If the **Working Directory** contains the mount folder for the virtual machine, remove it.
- 2 If you were using an advanced transport mode, such as **san** or **hotadd**, navigate to the `<system_drive>/windows/temp/vmware-system` directory.
- 3 If this directory contains any subdirectory named `<VM_UUID>-<VMmoref>`, delete it.

Here `<VM_UUID>` is the UUID of the mounted virtual machine and `VM_moref` is an internal reference that the ESXi or vCenter Server uses to see the virtual machine. To delete this folder, you may be required to set necessary permissions for the folder.

- 4 If you were using the **hotadd** transport mode, remove any disks of the target virtual machine—the virtual machine mounted for a backup—that were hotadded to the NetVault Backup Client Virtual Machine—the virtual machine where the Plug-in for VMware is running.

You can remove the disks from the vSphere Client for the virtual machine where the plug-in or NetVault Backup Client is running.

- 5 Go to the Snapshot Manager in vSphere Client, and remove the snapshot named **BKB_SNAP**, if it still exists.

Wait for a few minutes (approximately 2 to 3 minutes) to allow any memory cache to be cleared before performing this step.

If the virtual machine is powered on, an error message—“Unable to access file <unspecified filename> since it is locked”—might appear when you try to delete the snapshot, but the snapshot may thereafter disappear from the Snapshot Manager window. If a “Consolidate Helper-0” snapshot appears after you remove the snapshot, power off the virtual machine.

After removing the snapshot **BKB_SNAP**, VMware recommends that you create and delete a snapshot. You can create and delete snapshots from the Snapshot Manager window in vSphere Client. The server tries to consolidate the redo logs during this operation, and therefore, it may take a few minutes to complete. Delete the Consolidate Helper snapshot, if it exists.

- 6 If some unwanted redo logs are still present on the datastore, perform [Step 5](#) again.
- 7 Wait for a few minutes—approximately 2 to 3 minutes—to allow any memory cache to be cleared before performing this step.

Restoring image-level backups

- [About restoring image-level backups](#)
- [Restoring a full virtual machine or individual virtual disks](#)
- [Relocating a virtual machine to an alternate ESXi Server](#)
- [Renaming a virtual machine during restore](#)
- [Performing file-level restores from image-level backups](#)
- [Restoring virtual machine disk and configuration files](#)
- [Searching for files in savesets](#)
- [Viewing media list](#)

About restoring image-level backups

The image-level backups can be used to perform the following types of restores:

- **Recover a full virtual machine or specific virtual disks:** You can use image-level backups to recover a full virtual machine to a previous known state or to restore one or more virtual disks for a virtual machine. This method is useful when there is data loss due to hardware failure, data corruption, or accidental deletion of virtual machine disk files. The virtual machine can be restored to the same or an alternate VMware® ESXi™ Server Host.
- **Restore individual files and directories:** You can use image-level backups to restore individual files and folders. This method is useful when there is data loss due to user errors, data corruption, or accidental deletion of files. The individual files and directories can be restored to a specified directory on the NetVault Backup Client.

NOTE: To use an image-level backup for file-level restores, you must select the **Perform File Level Indexing** check box during backup.

File-level restores are supported on the following file systems:

- **Windows®:** NTFS
- **Linux® and UNIX®:** EXT2, EXT3, EXT4

The Plug-in for VMware 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.

The current version of the plug-in does not support Windows Server® 2012 ReFS (Resilient File System) and Striped Disks.

- **Restore virtual machine disk and configuration files:** You can use the image-level backups to restore the virtual machine disk and configuration files to a specified directory on the NetVault Backup Client. With these restored files, you can then recover a virtual machine with the same or modified settings using Virtual Infrastructure Client or any other utility that lets you create a virtual machine using existing `.vmdk` files.

Restoring a full virtual machine or individual virtual disks

The procedure for recovering a full virtual machine or individual virtual disks from an image-level backup includes the steps outlined in the following sections:

- [Prerequisites](#)
- [Restoring data](#)
- [Starting a virtual machine](#)

Prerequisites

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

- When restoring a virtual machine, the target virtual machine should not exist in the inventory. If the virtual machine that you want to restore still exists in the inventory, delete it. Alternatively, rename the virtual machine during restore.

When restoring individual virtual disks, the target virtual machine and virtual disks should exist in the inventory. If the virtual machine or virtual disks are unavailable, use the **Rename** option.

- Before you start the restore procedure on Windows® 2008, clear the read-only attribute for the target disk. If you fail to do so, the restore job completes successfully, but no data is restored; when you power on the restored virtual machine, it fails to boot.

To clear the read-only attribute for the target disk:

- 1 Start the **diskpart** utility, and list the disks:

```
Diskpart  
list disk
```

- 2 Select the target disk, and display the disk details:

```
Select disk <X>  
detail disk
```

- 3 If the read-only attribute is set to **Yes**, type:

```
attribute disk clear readonly
```

Restoring data

To restore a full virtual machine or individual virtual disks, use the following procedure.

- 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.

Table 11. Saveset status icons




Icon	Description
	Saveset is online (all segments are online).

Table 11. Saveset status icons

Icon	Description
	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.

- 2 To filter the saveset list, use the following filter options:

Table 12. Saveset filters

Filter	Description
Client	<p>Displays savesets created for particular clients. The default selection is Any.</p> <ol style="list-style-type: none"> 1 To use this filter, click the Client box. 2 In the Choose Client dialog box, select the clients. 3 To close the dialog box, click OK.
Plugin Type	<p>Displays savesets created using a particular plug-in. The default selection is Any.</p> <ol style="list-style-type: none"> 1 To use this filter, click the Plugin Type box. 2 In the list, select the plug-in.
Date	<p>Displays savesets created during a specified period. The default selection is Any.</p> <ol style="list-style-type: none"> 1 To use this filter, click the Date box. 2 In the list, select the option that you want to use. The available options are Last 24 hours, Last Week, Last Month, Last 6 Months, Last Year, and Any.
Job	<p>Displays savesets created for particular jobs. The default selection is Any.</p> <ol style="list-style-type: none"> 1 To use this filter, click the Job box. 2 In the Choose Job dialog box, select the jobs. 3 To close the dialog box, click OK.

- 3 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.

Backup tags indicate the type of image-level backup. The image-level backups use the following tags or backup-type identifiers:

- BACKUP IMAGE FULL
- BACKUP IMAGE INCREMENTAL
- BACKUP IMAGE DIFFERENTIAL

- 4 On the **Create Selection Set** page, select the images that you want to restore.

You can also open the images, and select the individual virtual disks that you want to restore.

NOTE: If you select an Incremental or a Differential Backup, the plug-in automatically restores all savesets in the backup sequence up to the selected saveset. The plug-in does not perform any unnecessary steps to read, send, or write disk sectors that have changed in a later backup. Each sector is restored only once.

- 5 Click **Edit Plugin Options**.
- 6 Depending on your VMware environment, configure the applicable options in the **Plug-in for VMware Restore Options** dialog box.
 - **Cluster setup:** In a cluster setup managed by a VMware® vCenter™ Server, configure the following options on the **Restore virtual machine under vCenter** tab.

Table 13. Options for restoring a full virtual machine in a cluster setup

Option	Description
Recover virtual machine under vCenter	To restore a full virtual machine or one or more virtual disks, select this option.
Allocate Resource Pool Path	<p>A resource pool represents the CPU and memory resources that are available to an ESXi Server Host or a cluster of hosts. These resources are made available to the individual virtual machines controlled by the host. Resource pools can be configured in terms of absolute minimum and maximum quantities or relative shares. To run a virtual machine, it must be allocated a resource pool.</p> <p>By default, a virtual machine is allocated to its original resource pool during restore. To allocate a virtual machine to an alternate resource pool, specify the target resource pool in this box. Use the following format to specify the resource pool:</p> <pre style="font-family: monospace;">/Pool-A/Pool-B/.../Pool<n></pre> <p>Here, <code>Pool<n></code> is the target resource pool, <code>Pool-A</code> is the child of the root resource pool, <code>Pool-B</code> is the child of <code>Pool-A</code>, and so on. This format allows you to specify a resource pool at any arbitrary depth within the hierarchy. To allocate the virtual machine to the root resource pool, type the “/” character.</p> <p>Note the following:</p> <ul style="list-style-type: none"> • If you do not specify any resource pool, the virtual machine is allocated to its original resource pool (as at the time of backup). • If you specify an invalid or inaccessible resource pool, the plug-in follows the path hierarchy backwards to find a valid resource pool. It allocates the virtual machine to the first valid ancestor pool in the path and logs an error message.
Alternate Datastore	<p>A datastore represents a storage location for virtual machine files. It can be a VMFS volume, directory on Network Attached Storage, or local file system path. A datastore is platform- and host-independent.</p> <p>To change the datastore for a virtual machine, specify the destination datastore name for the virtual machine. You can use this option when the original datastore is unavailable, inaccessible, or cannot hold the virtual machine due to lack of space. Verify that the name is correct and the datastore is accessible to the ESXi Server Host. When you configure an alternate datastore, all virtual disk and configuration files associated with a virtual machine are restored to a single datastore. Verify that the target datastore has sufficient disk space to hold the virtual machine files.</p> <p>Restore fails if the datastore is invalid or inaccessible or does not have sufficient space to store virtual machine files.</p> <ul style="list-style-type: none"> • Standalone ESXi Server setup: In a standalone ESXi Server setup, click the Restore virtual machine to standalone ESXi host tab, and configure the following options.

Table 14. Options for restoring a full virtual machine in a Standalone ESXi Server setup

Option	Description
Recover virtual machine to standalone ESXi host	To restore a full virtual machine or one or more virtual disks, select this option.
Alternate Datastore	To change the datastore for a virtual machine, specify the destination datastore name for the virtual machine. For more information about this option, see Alternate Datastore .

7 To save the settings, click **Save**, and then click **Next**.

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

Assign a descriptive name that lets you 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

9 In the **Target Client** list, verify that the client from which data was backed up is selected.

This client is selected by default. Do *not* change this setting.

10 Select or create the Schedule Set, Restore Source Set, and Advanced Options Set.

For more information about these sets, see the *Dell NetVault Backup Administrator's Guide*.

11 To submit the job for scheduling, click **Submit**.

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 *Dell NetVault Backup Administrator's Guide*.

Important notes

- If you select the virtual machine node, or select the virtual machine node and exclude one or more virtual disks, the plug-in creates a virtual machine and restores the contents of the selected disks.

When performing this type of restore, the target virtual machine should not exist in the inventory; otherwise, the restore job reports an error (“Found identical UUID”) and fails.

- If you select one or more virtual disks instead of the virtual machine, the plug-in overwrites the contents of the existing “.vmdk” file.

When restoring individual virtual disks, the target virtual machine and the selected virtual disks must exist in the inventory. If the virtual machine is removed from the inventory, the job reports an error (“Could not locate the VM”) and fails. The plug-in also reports an error if the target “.vmdk” file does not exist in the inventory.

- While restoring one or more virtual disks to an existing virtual machine, the plug-in tries to locate the virtual machine using its UUID. If the plug-in finds more than one virtual machine sharing the UUID, the following error message is logged and the job fails:

Failed to uniquely locate VM in server inventory because its uuid 'xxxx' is in use by one or more other VMs.

When this error occurs, you must change the UUID for the virtual machines in the inventory so that each virtual machine has a unique UUID. You can use the following procedure to locate virtual machines that use the same UUID.

To find virtual machines that share a UUID:

- 1 Open a browser window, and in the address bar, type the following URL.

```
https://<server_address>/mob/?moid=SearchIndex&method=findAllByUuid
```

Provide the login information for the server.









(<server_address> must be the address of the target vCenter or ESXi Server on which you want to restore the virtual disks.)

- 2 In the **uuid** box, type the UUID for the virtual machine that you want to restore.
- 3 In the **vmSearch** box, type **true**.
- 4 Leave the remaining options blank, and start the search.
- 5 The virtual machines that share the given UUID are displayed in the **val** list. Click the link for managed object reference of each virtual machine to view the name and other details.
- 6 Change the UUID so that each virtual machine has a unique UUID.

For more information about this procedure, see the relevant VMware documentation.

Restore selection tree icons

Table 15. Restore selection tree icons

Icon	Description
	Virtual machine
	Failed virtual machine (appears in a saveset for stopped job)
	Configuration file
	NVRAM file
	Virtual machine disk file
	Disk extent information file
	Log file
	Catalog file

Starting a virtual machine

After a virtual machine is restored, it is in powered off state. You must start it manually. When you start the virtual machine, the operating system might prompt to perform an error recovery or record the reason for the unclean shutdown. Usually, you can ignore this message and select the normal boot option.

Relocating a virtual machine to an alternate ESXi Server

- 1 Complete [Step 1](#) through [Step 4](#) in [Restoring data](#).
- 2 Click **Edit Plugin Options**.
- 3 Depending on your VMware environment, configure the applicable options in the **Plug-in for VMware Restore Options** dialog box.
 - **Cluster setup:** To relocate a virtual machine in a cluster setup, configure the following options on the **Restore virtual machine** under **vCenter** tab.

Table 16. Relocating a virtual machine in a cluster setup

Option	Description
Recover virtual machine under vCenter	To restore a full virtual machine or one or more virtual disks, select this option.
Alternate Datastore	Specify the destination datastore name for the virtual machine. For more information about this option, see Alternate Datastore . This option is required when relocating a virtual machine. If you do not specify this job, the restore job fails.

- **Standalone ESXi Server setup:** To relocate a virtual machine to a standalone ESXi Server, click the **Restore virtual machine to standalone ESXi** host tab, and configure the following options.

Table 17. Relocating a virtual machine to a standalone ESXi Server

Option	Description
Recover virtual machine under vCenter	To restore a full virtual machine or one or more virtual disks, select this option.
Alternate Datastore	Specify the destination datastore name for the virtual machine. For more information about this option, see Alternate Datastore . This option is required when relocating a virtual machine. If you do not specify this job, the restore job fails.
Host	Type the IP address or DNS name of the alternate ESXi Server host.
User Name	Specify a user account that can be used to log on to the ESXi Server. The user account must have privileges to register or create virtual machines
Password	Type the password for the user account.

- 4 Complete [Step 5](#) through [Step 11](#) in [Restoring data](#).

NOTE: A feature introduced in vSphere 5 prevents you from restoring a virtual machine directly to an ESXi 5 host when the host is managed by a vCenter Server. To restore a virtual machine directly to the host, you must first disassociate the host from the vCenter Server.

To disassociate an ESXi Server from a vCenter Server:

- 1 From the vSphere Client, connect directly to the ESXi 5 host.
- 2 In the Inventory panel, select the host.
- 3 In the panel on the right, click **Summary**.
- 4 Select the **Disassociate host from vCenter Server** check box.

For more information, see the relevant vSphere documentation.

Renaming a virtual machine during restore

To restore a virtual machine or its individual disks without overwriting the existing copy, you can rename the virtual machine during restore. The plug-in creates a virtual machine and restores the contents of the selected disks.

- 1 Complete [Step 1](#) through [Step 4](#) in [Restoring data](#).
- 2 Select the target virtual machine, and in the **Actions** list, click **Rename**.

NOTE: Only backups created with Plug-in for VMware 2.6 and later support the **Rename** feature.


- In the **Restore Rename** dialog box, select the **Rename To** check box, and type a new name for the virtual machine in the associated box.

The name can contain a maximum of 80 characters. It cannot include any special characters (such as, / \ ? * : @ > < | " ? &). If you specify more than 80 characters or use any special character in the virtual machine name, the job reports an error and fails.

- To close the dialog box, click **OK**.

After you close the dialog box, the plug-in updates the corresponding node in the selections tree to display the new name for the virtual machine.

- Complete [Step 5](#) through [Step 11](#) in [Restoring data](#).

 **NOTE:** When you rename a virtual machine, the plug-in renames all disk and configuration files except the `vmware.log` files.

- When you start the renamed virtual machine and a dialog box is displayed that asks whether you copied the virtual machine or moved it, select the **"I copied it"** option, and close the dialog box.




Performing file-level restores from image-level backups

- 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.

Table 18. Saveset status 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.

- To filter the saveset list, use the following filter options:

Table 19. Saveset filters

Filter	Description
Client	Displays savesets created for particular clients. The default selection is Any . <ol style="list-style-type: none"> To use this filter, click the Client box. In the Choose Client dialog box, select the clients. To close the dialog box, click OK.
Plugin Type	Displays savesets created using a particular plug-in. The default selection is Any . <ol style="list-style-type: none"> To use this filter, click the Plugin Type box. In the list, select the plug-in.

Table 19. Saveset filters

Filter	Description
Date	<p>Displays savesets created during a specified period. The default selection is Any.</p> <ol style="list-style-type: none"> To use this filter, click the Date box. In the list, select the option that you want to use. The available options are Last 24 hours, Last Week, Last Month, Last 6 Months, Last Year, and Any.
Job	<p>Displays savesets created for particular jobs. The default selection is Any.</p> <ol style="list-style-type: none"> To use this filter, click the Job box. In the Choose Job dialog box, select the jobs. To close the dialog box, click OK.

- 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.

- On the **Create Selection Set** page, open the applicable container nodes—that is, the vCenter, Datacenter, ESXi Host Cluster, and ESXi Server—to list the virtual machine images stored in the saveset.
- Open the image from which you want to restore the files.
- Open the **Volumes** node, and then open the disk partitions to display the directory tree.
- Select the files and directories that you want to restore.
- Click **Edit Plugin Options**.
- In the **Plug-in for VMware Restore Options** dialog box, click the **Restore Files** tab, and configure the following options:

Table 20. Options for file-level restores from image-level backups

Option	Description
Files Restore	<p>To perform file-level restores from image-level backups, select this option. After selecting the individual files and directories under the Volumes node, if you do not select this option, the job fails.</p>
Target Directory	<p>Type the full path to the directory to which you want to restore the virtual machine disk and configuration files. The target directory must be local to the NetVault Backup Client on which the Plug-in for VMware is running. Currently, mapped network drives, mounted network shares, and UNC paths are not supported.</p> <p>The plug-in does not check if you have specified a local or remote directory. If plug-in cannot access the specified path, the job fails.</p>

- To save the settings, click **Save**, and then click **Next**.

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

Assign a descriptive name that lets you 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.

- In the **Target Client** list, verify that the client from which data was backed up is selected.

This client is selected by default. Do *not* change this setting.

- Select or create the Schedule Set, Restore Source Set, and Advanced Options Set.











For more information about these sets, see the *Dell NetVault Backup Administrator's Guide*.

14 To submit the job for scheduling, click **Submit**.

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 *Dell NetVault Backup Administrator's Guide*.

Restore selection tree icons

Table 21. Restore selection tree icons

Icon	Description
	Virtual machine
	Failed virtual machine (appears in a saveset for stopped job)
	Volume
	Partition
	Configuration file
	NVRAM file
	Virtual machine disk file
	Disk extent information file
	Log file
	Catalog file

Restoring virtual machine disk and configuration files

The procedure for restoring virtual machine disk and configuration files from an image-level backup includes the steps outlined in the following sections:

- [Restoring data](#)
- [Recovering a virtual machine from the restored files](#)

Restoring data




To restore virtual machine disk and configuration files, use the following procedure.

- 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.

Table 22. Saveset status 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.

- To filter the saveset list, use the following filter options:

Table 23. Saveset filters

Filter	Description
Client	<p>Displays savesets created for particular clients. The default selection is Any.</p> <ol style="list-style-type: none"> To use this filter, click the Client box. In the Choose Client dialog box, select the clients. To close the dialog box, click OK.
Plugin Type	<p>Displays savesets created using a particular plug-in. The default selection is Any.</p> <ol style="list-style-type: none"> To use this filter, click the Plugin Type box. In the list, select the plug-in.
Date	<p>Displays savesets created during a specified period. The default selection is Any.</p> <ol style="list-style-type: none"> To use this filter, click the Date box. In the list, select the option that you want to use. The available options are Last 24 hours, Last Week, Last Month, Last 6 Months, Last Year, and Any.
Job	<p>Displays savesets created for particular jobs. The default selection is Any.</p> <ol style="list-style-type: none"> To use this filter, click the Job box. In the Choose Job dialog box, select the jobs. To close the dialog box, click OK.

- 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.

- On the **Create Selection Set** page, select the image from which you want to restore the disk and configuration files.
- Click **Edit Plugin Options**.
- In the **Plug-in for VMware Restore Options** dialog box, click the **Restore Files** tab, and configure the following options:

Table 24. Options for restoring virtual machine disk and configuration files from image-level backups

Option	Description
Files Restore	To perform file-level restores from image-level backups, select this option. NOTE: After selecting the individual files and directories under the Volumes node, if you do not select this option, the job fails.
Target Directory	Type the full path to the directory to which you want to restore the virtual machine disk and configuration files. The target directory must be local to the NetVault Backup Client on which the Plug-in for VMware is running. Currently, mapped network drives, mounted network shares, and UNC paths are not supported. The plug-in does not check if you have specified a local or remote directory. If plug-in cannot access the specified path, the job fails.
Split Restored Virtual Disks	This option can be used when restoring virtual disk files to the local file system. If the check box is selected, the plug-in splits the “.vmdk” file into 2GiB files. If you clear the check box, the plug-in restores the “.vmdk” file as a single disk file. This option is only available to backups created using Plug-in for VMware 10.0.5 and later. To force this behavior when restoring a backup created using Plug-in for VMware 10.0.1 or an older version, configure this option in the vmware.cfg file: <ol style="list-style-type: none"> 1 Open the vmware.cfg file in any text editor. You can find this file in <NetVault Backup home>\config on Windows® and <NetVault Backup home>/config on Linux®. 2 Add the following lines: [Custom:RestoreSplitVMDK] Value=TRUE 3 Save the file. The setting for this option in the Plug-in for VMware Restore Options dialog box overrides the setting in the vmware.cfg file.

7 To save the settings, click **Save**, and then click **Next**.

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

Assign a descriptive name that lets you 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.

9 In the **Target Client** list, verify that the client from which data was backed up is selected.

This client is selected by default. Do *not* change this setting.

10 Select or create the Schedule Set, Restore Source Set, and Advanced Options Set.









For more information about these sets, see the *Dell NetVault Backup Administrator's Guide*.

11 To submit the job for scheduling, click **Submit**.

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 *Dell NetVault Backup Administrator's Guide*.

Restore selection tree icons

Table 25. Restore selection tree icons

Icon	Description
	Virtual machine
	Failed virtual machine (appears in a saveset for stopped job)
	Configuration file
	NVRAM file
	Virtual machine disk file
	Disk extent information file
	Log file
	Catalog file

Recovering a virtual machine from the restored files

After restoring the virtual machine disk and configuration files, you can use the VMware vCenter Converter Standalone Client to create a virtual machine using the restored “.vmx” and “.vmdk” files. After a VADP Recovery, you must modify the “.vmx” file so that it matches the disk paths of the restored “.vmdk” files. For more information, see the relevant VMware vCenter Converter Standalone Client documentation.

Recovering a virtual machine from a VCB backup

To restore the virtual machine files to the destination VMware Server, start the VMware vCenter Converter Standalone Client Wizard. Follow the instructions to complete the virtual machine recovery. When the prompt “Select source type” prompt appears, select the option “Backup image or third-party virtual machine.” Also, set the “Virtual machine file” option to the restored “.vmx” file.

Recovering a virtual machine from a VADP backup using vCenter Converter

To recover the virtual machine, you must edit the “.vmx” file, and change the datastore path to reflect the restored “.vmdk” files.

To change the datastore path:

- 1 Navigate to the target directory to which the virtual machine files were restored and note the filenames of the virtual disks.

The following example shows some files generated from Virtual Machine Files Only Restore of a VADP backup:

```
MyVirtualMachine.vmx
scsi0-0-MyVirtualMachine.vmdk
scsi0-0-MyVirtualMachine-s001.vmdk
VirtualMachineConfigInfo
```

- 2 Open the “.vmx” file.
- 3 Modify the relevant entries for the virtual disks.

For example, modify the following entry:

```
scsi0:0.fileName = "MyVirtualMachine-000001.vmdk"
```

It must match the “.vmdk” file generated during restore.

```
scsi0:0.fileName = "scsi0-0-MyVirtualMachine.vmdk"
```

- 4 After modifying the “.vmx” file, start the VMware vCenter Converter Standalone Client Wizard to create the virtual machine.
- 5 Follow the instructions to complete the virtual machine recovery.
- 6 When the “Select source type” prompt appears, select “VMware Workstation or other VMware virtual machine.”
- 7 Set the “Virtual machine file” option to the restored “.vmx” file.

① | **NOTE:** In the scenario described in the article <http://kb.vmware.com/kb/1019286>, the “.vmx” file is not backed up, but the “.vmdk” files are available.

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 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 in the **Search String** box, select this check box.
 - **Case sensitive:** To perform a case-sensitive search, select this check box.
- 3 Click **Search**.

On the **Search Results** page, you can view 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 **Next**.
- 6 Complete [Step 5](#) through [Step 11](#) in [Restoring data](#).

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.

- 1 On the **Create Restore Job – Choose Saveset** page, select the applicable saveset.
- 2 In the **Saveset Information** area, click **Media List**.

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.
- 3 To close the dialog box, click **Close**.

Restoring file-level backups

- [About restoring file-level backups](#)
- [Restoring file-level backups using shared network drives](#)
- [Restoring file-level backups using the native Plug-in for FileSystem](#)
- [Searching for files in savesets](#)
- [Viewing media list](#)

About restoring file-level backups

The Plug-in for VMware provides the following methods for restoring file-level backups.

- **Restores using shared network drives:** This method does not require an instance of NetVault Backup Client running inside the virtual machine. Restores are performed to the shared network drives on the virtual machines, using the Plug-in for FileSystem installed on the plug-in host. Since the Plug-in for FileSystem is automatically installed on the NetVault Backup Client, no additional software is required for this type of restore.
- **Restores using native Plug-in for FileSystem:** This method is suitable if an instance of NetVault Backup Client is already running inside the virtual machine. The virtual machine is treated as any other physical NetVault Backup Client for the file system restores.

Restoring file-level backups using shared network drives

The procedure for restoring file-level backups to a shared folder on the virtual machine includes the steps outlined in the following sections:

- [Setting up a network share](#)
- [Restoring data](#)

Setting up a network share

To create a network share between the NetVault Backup Client and the target virtual machine, use the following procedure.

- 1 On the target virtual machine, configure the sharing properties for a folder.
- 2 On the NetVault Backup Server, start the Backup Job Wizard, and click **Create New** next to the **Selections** list.
- 3 On the **NetVault Backup Selections** page, open the NetVault Backup Client on which the plug-in is installed, and then open **File System**.

- 4 Select **Network Shares**, and in the **Actions** list, click **Add Network Share**.
- 5 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>`
- 6 Click **Add** to add the network share and close the dialog box.
- 7 Select **Network Shares**, and in the **Actions** list, click **Connect As**.
- 8 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.
 - ① **NOTE:** If you configure a Non-Domain Administrator account for backups, the file and directory permissions must be set manually after a restore.
 - **Password:** Type the password for the user account.
- 9 To save the connection details and close the dialog box, click **OK**.

Restoring data




To restore file-level backups, use the following procedure.

- 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.

Table 26. Saveset status 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.

- 2 To filter the saveset list, use the following filter options:

Table 27. Saveset filters

Filter	Description
Client	Displays savesets created for particular clients. The default selection is Any . <ol style="list-style-type: none"> 1 To use this filter, click the Client box. 2 In the Choose Client dialog box, select the clients. 3 To close the dialog box, click OK.
Plugin Type	Displays savesets created using a particular plug-in. The default selection is Any . <ol style="list-style-type: none"> 1 To use this filter, click the Plugin Type box. 2 In the list, select the plug-in.

Table 27. Saveset filters

Filter	Description
Date	<p>Displays savesets created during a specified period. The default selection is Any.</p> <ol style="list-style-type: none"> To use this filter, click the Date box. In the list, select the option that you want to use. <p>The available options are Last 24 hours, Last Week, Last Month, Last 6 Months, Last Year, and Any.</p>
Job	<p>Displays savesets created for particular jobs. The default selection is Any.</p> <ol style="list-style-type: none"> To use this filter, click the Job box. In the Choose Job dialog box, select the jobs. To close the dialog box, click OK.

- 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.

- On the **Create Selection Set** page, select the files and directories that you want to restore.

NOTE: For a regular Incremental or Differential Backup, the selection tree displays all files that were included in the backup sequence—the initial Full Backup and subsequent Incremental or Differential Backups. For a Dump Type Incremental or Differential Backup, the plug-in only displays those files that were backed up in the selected saveset.

- Select the file or directory that you want to restore, and in the **Actions** list, click **Rename**.

- In the **Relocate** box, type the network share path.

NOTE: If you do not perform [Step 5](#) and [Step 6](#), the data is restored to the NetVault Backup Client and not to the shared folder on the virtual machine.

- Click **Edit Plugin Options**, and configure the following settings:

Table 28. Plug-in for FileSystem restore options

Option	Description
Overwrite Newer Files	<p>By default, the plug-in overwrites existing files on the restore destination with the backed-up version in the saveset.</p> <p>If you do not want to overwrite the existing files, clear this check box.</p>
Reset File Timestamps	<p>By default, timestamps for restored files are reset to what is recorded in the backup saveset.</p> <p>To restore the files with current timestamp, clear the check box.</p>
Reset Directory Timestamps	<p>By default, directories are restored with current timestamp.</p> <p>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.</p>
Restore Files Modified During Backup	<p>To restore the files that were marked “in flux” during backup, select this check box.</p>

Table 28. Plug-in for FileSystem 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, 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 the <i>Dell NetVault Backup Plug-in for FileSystem User’s Guide</i> .

8 To save the settings, click **Save**, and then click **Next**.

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

Assign a descriptive name that lets you 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

10 In the **Target Client** list, verify that the client from which data was backed up is selected.

This client is selected by default. Do *not* change this setting.

11 Select or create the Schedule Set, Restore Source Set, and Advanced Options Set.

For more information about these sets, see the *Dell NetVault Backup Administrator’s Guide*.

12 To submit the job for scheduling, click **Submit**.

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 *Dell NetVault Backup Administrator’s Guide*.

① | **NOTE:** The plug-in restores data to the shared network folder on the virtual machine. You must manually move the restored files to the final target destination.

Restoring file-level backups using the native Plug-in for FileSystem

The procedure for restoring the file-level backups using the native Plug-in for FileSystem includes the steps outlined in the following steps:

- [Prerequisites](#)
- [Restoring data](#)

Prerequisites

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

- Install the NetVault Backup Client software inside the target virtual machine. For more information, see the *Dell NetVault Backup Installation Guide*.
- Add the client to NetVault Backup Server. For more information, see the *Dell NetVault Backup Administrator’s Guide*.

Restoring data




To select data for a restore, use the following procedure.

- 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.

Table 29. Saveset status 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.

- 2 To filter the saveset list, use the following filter options:

Table 30. Saveset filters

Filter	Description
Client	Displays savesets created for particular clients. The default selection is Any . <ol style="list-style-type: none">1 To use this filter, click the Client box.2 In the Choose Client dialog box, select the clients.3 To close the dialog box, click OK.
Plugin Type	Displays savesets created using a particular plug-in. The default selection is Any . <ol style="list-style-type: none">1 To use this filter, click the Plugin Type box.2 In the list, select the plug-in.
Date	Displays savesets created during a specified period. The default selection is Any . <ol style="list-style-type: none">1 To use this filter, click the Date box.2 In the list, select the option that you want to use. The available options are Last 24 hours, Last Week, Last Month, Last 6 Months, Last Year, and Any.
Job	Displays savesets created for particular jobs. The default selection is Any . <ol style="list-style-type: none">1 To use this filter, click the Job box.2 In the Choose Job dialog box, select the jobs.3 To close the dialog box, click OK.

- 3 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.

- 4 On the **Create Selection Set** page, select the files and directories that you want to restore.

NOTE: For a regular Incremental or Differential Backup, the selection tree displays all files that were included in the backup sequence—the initial Full Backup and subsequent Incremental or Differential Backups. For a Dump Type Incremental or Differential Backup, the plug-in only displays those files that were backed up in the selected saveset.

- 5 Click **Edit Plugin Options**, and configure the applicable settings.

For more information, see [Plug-in for FileSystem restore options](#).

- 6 To save the settings, click **Save**, and then click **Next**.

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

Assign a descriptive name that lets you 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.

- 8 In the **Target Client** list, select the target virtual machine.

IMPORTANT: If you do not select the virtual machine in the Target Client list, the data is restored to the client on which the plug-in is installed.

- 9 Select or create the Schedule Set, Restore Source Set, and Advanced Options Set.

For more information about these sets, see the *Dell NetVault Backup Administrator's Guide*.

- 10 To submit the job for scheduling, click **Submit**.

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 *Dell NetVault Backup 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.

- 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 in the **Search String** box, select this check box.
- **Case sensitive:** To perform a case-sensitive search, select this check box.

- 3 Click **Search**.

On the **Search Results** page, you can view 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 **Next**.

- 6 Complete [Step 5](#) through [Step 12](#) in [Restoring data](#).

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.

- 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, 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 To close the dialog box, click **Close**.

Troubleshooting

- [Common errors](#)
- [Diagnosing virtual machine issues](#)
- [SOAP messages](#)
- [Generating VDDK logs](#)

Common errors

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

- [NetVault Backup Service fails to start on Windows](#)
- [NetVault Backup Service fails to start after the machine is restarted](#)
- [NetVault Backup Service starts, but stops immediately on Linux](#)
- [Linux Hybrid version of the plug-in does not run on 64-bit OS](#)
- [Crash occurs during a backup](#)
- [CBT-based image-level backup fails](#)
- [SAN transport reports errors when multiple backups access the same datastore](#)
- [Default SCSI controller type on RHEL 7 and RHEL 6](#)
- [File-level backups do not back up data for a reparese point](#)
- [Restore job fails when using SAN transport](#)
- [Restore job fails to add a virtual machine to the server inventory](#)
- [Backup jobs of virtual machines protected by vSphere FT fail intermittently](#)

NetVault Backup Service fails to start on Windows

Description

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

NetVault Backup Service fails to start after the machine is restarted

Description

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

NetVault Backup Service starts, but stops immediately on Linux

Description

On a Linux®-based machine, the NetVault Backup 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.

Linux Hybrid version of the plug-in does not run on 64-bit OS

Description

The Linux® Hybrid version of the plug-in does not run on pure 64-bit operating system.

Symptom

The plug-in fails to open and displays the message "No Error."

Solution

Pure 64-bit operating systems do not install some common 32-bit libraries the plug-in requires.

To correct this problem, do the following:

- 1 Start a terminal session, and go to `/usr/netvault/bin`.
- 2 To list the missing libraries, run the following command:

```
ldd nvvmware | egrep -i "missing|not found"
```

3 Use the applicable package manager to install the missing 32-bit libraries.

① | **NOTE:** Some libraries are dynamically loaded from the path `$NV_HOME/dynlib/vddk/`. If these libraries are listed as missing, no action is required.

Crash occurs during a backup

Description

During a backup, the plug-in reports an error and crashes.

Symptom

The plug-in crashes with the following error in the NetVault Backup logs:

Backing up data of VM disk [XXXXXXX]xxx/xxxxxxxxx.vmdk' to media : "Job manager lost message channel unexpectedly."

The trace logs include the following message:

Cannot open library: libexpat.so.0: cannot open shared object file in vixDiskLib-16642.log.

Solution

To correct this problem, navigate to the `/usr/lib` directory, and create the following symbolic link:

```
ln -s /lib/libexpat.so.1.5.2 libexpat.so.0
```

CBT-based image-level backup fails

Description

An image-level backup for which CBT has been enabled reports an error and fails.

Symptom

The job fails with the following error in the NetVault Backup logs:

Failed to get changed disk areas.

Solution

This issue occurs when a snapshot exists on a virtual machine when CBT is enabled. For snapshots that exist before enabling CBT, the `changeld` parameter is not set. So, when the `QueryChangedDiskAreas` API is invoked, it returns an error. For more information, see <http://kb.vmware.com/kb/1033816>.

To correct this problem, do either of the following:

- Ensure that no snapshots exist on the virtual machine before selecting the **Enable Changed Block Tracking for VM** check box.
- or —
- Use the **Reset Change Block Tracking** option. This option reconfigures CBT on a virtual machine and allows you to run CBT-enabled backup that failed previously. For more information about this option, see [Resetting CBT on a virtual machine](#).

SAN transport reports errors when multiple backups access the same datastore

Description

A Linux[®]-based client may fail to read the contents of a virtual disk using the SAN transport mode if multiple backup processes are accessing the same datastore.

Symptom

During a backup using the SAN transport mode, the following errors are reported:

San transport error: I/O Operation failed.

Error: One of the parameters supplied is invalid.

Solution

These errors are caused by a fault in the Linux kernel when handling SCSI reservation conflicts.

The errors are encountered on the following systems:

- Red Hat[®] Enterprise Linux[®] (RHEL) 6.1 and earlier versions
- SUSE[®] Linux Enterprise Server (SLES[®]) 11.1 and earlier versions

To correct this issue, upgrade to kernel version 2.6.32-220 or later on RHEL and kernel version 3.0.13 or later on SLES.

Default SCSI controller type on RHEL 7 and RHEL 6

Description

When you create a virtual machine with RHEL 7 or RHEL 6 as the Guest OS, the SCSI controller type is set to **VMware Paravirtual** by default. The HotAdd transport mode does not work if either the backup proxy or the backup target uses Paravirtual SCSI controller.

Symptom

N/A

Solution

To use HotAdd transport mode, set the SCSI controller type to **LSI Logic Parallel** or **LSI Logic SAS** when you create the virtual machine. For more information about configuring this setting, see the relevant VMware documentation.

Dell does not recommend changing the SCSI controller type for an existing virtual machine without testing to ensure that it operates normally with the new controller type.

File-level backups do not back up data for a reparse point

Description

The file-level backup method does not back up data through a reparse point.

Symptom

When a disk mounted as a reparse point on a Guest OS is included in a file-level backup, the backup completes without any errors or warnings. However, the plug-in does not back up any data for a reparse point. It is listed as an empty node on the **Create Selection Set** page.

Solution

The VMware API does not support this type of backup.

Workaround:

- Install NetVault Backup Client on the Guest OS, and perform a backup using the Plug-in for FileSystem to include the reparse point directory.
—or—
- Mount the disk as drive letter instead of a reparse point. Then, create a backup job to include the new drive letter.

Restore job fails when using SAN transport

Description

Restore job fails.

Symptom

A restore job fails with the following error:

Log Message:

Failed to restore VM disk data.

Log Context:

Failed to open vmdk VDDK error encountered: code 13

Retrieving error message text VDDK error message is 'You do not have access rights to this file'

Disk data restore failed for '[datastore]VirtualMachineName/ VirtualDiskName.vmdk'

Solution

The NetVault Backup Client running the Plug-in for VMware may not have access to the SAN. To verify, try using the NBD transport mode. If restore completes successfully with this mode, it indicates SAN connectivity issues.

To use the SAN transport mode, the NetVault Backup Client must have SAN connectivity to the VMFS LUN. Ensure that the LUN is accessible and the SAN is configured correctly.

Restore job fails to add a virtual machine to the server inventory

Description

Restore job fails.

Symptom

A restore job fails with the following error:

Log message:

Failed to add VM to target server's inventory.

Log context:

Fault string is 'Invalid configuration for device '3'.

Solution

When this error occurs, enable trace and run the restore job again. When the job completes, open the file named `vmw<nnn>_soaprecv.log`, and check if it contains an XML message `<WaitForUpdatesExResponse>` with contents similar to the following:

```
<val xsi:type="LocalizedMethodFault">
<fault xsi:type="InvalidDeviceSpec">
<property>deviceChange[3].device.backing.fileName
</property>
<deviceIndex>3</deviceIndex>
</fault>
<localizedMessage>Invalid configuration for device &apos;3&apos;.</localizedMessage>
</val>
```

If so, it indicates that the backup contains an invalid configuration for a CD or DVD drive.

To correct this problem:

- 1 Open the file `vmware.cfg` in a text editor.

You can find this file in `<NetVault Backup home>\config` on Windows® and `<NetVault Backup home>/config` on Linux®.

- 2 Set the value for the following entry to **True**:

```
[Custom:ReconfigureVirtualCdromDevices]
Value=True
```

(The default value for this entry is **False**.)

- 3 Save the file.

- 4 Run the restore job again.

The job completes successfully. If a CD or DVD drive was set to “Host Device” or “Datastore ISO File” type during the original backup, it is reconfigured as “Client Device” type. A warning message in the logs provides a list of reconfigured devices. The log context includes their original configuration. To view the log context, click the **More Info** button.

- 5 After restoring the virtual machine, reset `[Custom:ReconfigureVirtual CdromDevices]` to **False**.

Backup jobs of virtual machines protected by vSphere FT fail intermittently

Description

When backing up virtual machines that are protected using VMware vSphere® Fault Tolerance (vSphere FT), the backup job fails.

Symptom

The job fails and reports the following: Cannot backup virtual machine while legacy fault tolerance is enabled.

Solution

 **WARNING:** Before using the plug-in to back up virtual machines that are protected using vSphere FT, make sure that your environment uses VMware ESXi 6.0 build number 4192238 or later.

Your virtual machine is protected by an earlier version of fault tolerance that does not support backup snapshots.

To back up and restore virtual machines that are using legacy fault tolerance, reconfigure the virtual machines to use a newer version of vSphere FT. Using the vSphere Web Client, turn off fault tolerance for the virtual

machine, and then turn it on again. During the reconfiguration process, FT protection is *not* available for the virtual machine that you are reconfiguring.

Diagnosing virtual machine issues

The **Diagnose Virtual Machine** method includes some predefined tests, which are used to identify issues in virtual machine setup that can cause failures when the plug-in performs the following activities:

- Access the VMDK files
- Check CBT status and capabilities
- Create quiesced snapshots

The plug-in reports the problems detected during these tests, and provides suggestions to correct the problems. You can run these tests at any time.

To run diagnostic tests:

- 1 Start the Backup Job Wizard, and click **Create New** next to the **Selections** list.
- 2 Open the NetVault Backup Client on which plug-in is installed, and then open **VMware Plugin**.
- 3 Open the VMware® ESXi™ or VMware® vCenter™ Server and the other applicable container nodes—for example, Datacenter, Cluster, Resource Pool, and other nodes—to display the target virtual machine.
- 4 Select the virtual machine, and in the **Actions** list, click **Diagnose Virtual Machine**.

The plug-in performs the following tests:

- Create a snapshot.
- ① **NOTE:** When performing diagnostic tests, the plug-in tries to create a quiesced snapshot. This operation may take some time to complete depending on how busy the virtual machine is.
- Open and read VMDK files using the configured primary transport mode.
 - Check the virtual hardware version.
 - Check the current CBT setting.
 - Check if VMware Tools are installed and running.
 - Check if any snapshots exist on the virtual machine.
- 5 View the results in the **Diagnostic Results** dialog box:
 - **Results:** This tab shows the results of the diagnostic tests (passed or failed).
 - **Recommendations:** This tab includes suggestions to correct the problems detected during the tests.
 - **Virtual Machine:** This tab shows general information about the virtual machine.
 - 6 To close the dialog box, click **Close**.

SOAP messages

When you enable tracing, the Plug-in for VMware generates the following two files along with the other trace logs:

- **vmw(nnn)_soapsent.log**
- **vmw(nnn)_soaprecv.log**

These files contain communication details between the plug-in and vSphere Web Service running on the VMware vCenter or ESXi Servers.

When you send the trace files to Dell Technical Support, ensure that these two files are included.

Generating VDDK logs

The Plug-in for VMware uses the VMware VDDK API for the following operations:

- Image-level backups and restores
- File-level browse and backups

To diagnose or troubleshoot errors encountered during these operations, Dell Technical Support may request you to send the VDDK logs. You can use the following procedure to enable VDDK logging when these operations are performed.

To generate VDDK logs:

- 1 In the navigation pane, click **Create Backup Job**, and then click **Create New** next to the **Selections** list.
- 2 Open the NetVault Backup Client on which the plug-in is installed.
- 3 Select **VMware Plugin**, and in the **Actions** list, click **Configure**.

NOTE: You can also configure the default settings from the **Change Settings** page:

- 1 In the Navigation pane, click **Change Settings**.
 - 2 If the plug-in is installed on the NetVault Backup Server, click **Server Settings**.
— or —
If the plug-in is installed on a NetVault Backup Client, click **Client Settings**; in the clients table, select the client, and click **Next**.
 - 3 Under **Plugins**, click **Plugin Options**.
- 4 Under **Troubleshooting**, select the **Enable VDDK Logging** check box.
 - 5 To save the settings, click **OK** or **Apply**.

The plug-in generates the VDDK log files entitled **vixDiskLib*.log** in the following directories:

- **Windows**[®]: <System Drive>\Windows\Temp\vmware-SYSTEM
- **Linux**[®]: /tmp/vmware-root

The log files are generated when you perform image-level backup or restore or file-level backup or browse operations.

Dell listens to customers and delivers worldwide innovative technology, business solutions, and services they trust and value. For more information, visit <http://www.software.dell.com>.

Contacting Dell

For sales or other inquiries, visit <http://software.dell.com/company/contact-us.aspx> or call +1-949-754-8000.

Technical support resources

Technical support is available to customers who have purchased Dell software with a valid maintenance contract and to customers who have trial versions. To access the Support Portal, go to <https://support.software.dell.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. In addition, the Support Portal provides direct access to product support engineers through an online Service Request system.

The Support Portal enables you to:

- Create, update, and manage Service Requests (cases).
- View Knowledge Base articles.
- Obtain product notifications.
- Download software. For trial software, go to <http://software.dell.com/trials>.
- View how-to videos.
- Engage in community discussions.
- Chat with a support engineer.