

Dell™ NetVault™ Backup
Plug-in for VMware® 10.0.5
User's Guide



© 2015 Dell Inc.
ALL RIGHTS RESERVED.

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

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

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

Dell Inc.
Attn: LEGAL Dept
5 Polaris Way
Aliso Viejo, CA 92656

Refer to our web site (software.dell.com) for regional and international office information.


Patents


This product is protected by U.S. Patents # 7,814,260; 7,913,043; 7,979,650; 8,086,782; 8,145,864; 8,171,247; 8,255,654; 8,271,755; 8,311,985; and 8,452,731. Protected by Japanese, E.U., French, and UK patents 1615131 and 05250687.0, and German patent DE602004002858. Additional patents pending. For more information, go to <http://software.dell.com/legal/patents.aspx>.


Trademarks

Dell, the Dell logo, and NetVault are trademarks of Dell Inc. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. VMware, VMware vCenter, VMware ESX, 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. SQL Server, SharePoint, and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims any proprietary interest in the marks and names of others.

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 - April 2015
Software Version - 10.0.5
VMW-101-10.0.5-EN-01

Contents

| | |
|--|----|
| Introduction | 5 |
| About Dell™ NetVault™ Backup Plug-in for VMware® | 5 |
| Key benefits | 5 |
| Feature summary | 6 |
| Target audience | 6 |
| Recommended additional reading | 7 |
| Installing the plug-in | 8 |
| About plug-in editions | 8 |
| Deploying the plug-in | 8 |
| Prerequisites | 9 |
| Installing the plug-in | 11 |
| Removing the plug-in | 11 |
| Configuring the plug-in | 12 |
| Adding a server | 12 |
| Reconfiguring a server | 13 |
| Removing a server | 13 |
| Configuring default settings | 14 |
| Configuring CBT | 16 |
| Enabling CBT on a virtual machine | 16 |
| Disabling CBT on a virtual machine | 17 |
| Manually disabling CBT on a virtual machine | 17 |
| Enabling virtual machine locking | 18 |
| Understanding backup consistency states | 18 |
| Defining a backup strategy | 20 |
| About backup methods and types | 20 |
| Image-level backups | 20 |
| File-level backups | 22 |
| Defining a backup and recovery strategy | 23 |
| Using the image-level backup method | 24 |
| Performing image-level backups | 24 |
| Additional notes | 29 |
| Backup selection tree icons | 31 |
| Resetting CBT on a virtual machine | 32 |
| Monitoring job progress | 33 |
| Using the file-level backup method | 34 |
| Performing file-level backups | 34 |
| Backup selection tree icons | 36 |
| Additional notes | 37 |
| Manually removing the snapshot and mount folder | 37 |

| | |
|--|-----------|
| Restoring image-level backups | 39 |
| About restoring image-level backups | 39 |
| Restoring a full virtual machine or individual virtual disks | 40 |
| Prerequisites | 40 |
| Restoring data | 40 |
| Starting a virtual machine | 43 |
| Using other restore procedures | 44 |
| Relocating a virtual machine to an alternate ESX Server host | 44 |
| Renaming a virtual machine during restore | 44 |
| Performing file-level restores from image-level backups | 45 |
| Restoring virtual machine disk and configuration files | 47 |
| Restoring data | 47 |
| Recovering a virtual machine from the restored files | 49 |
| Using additional features available on the Choose Saveset page | 50 |
| Filtering the saveset list | 50 |
| Searching for files in savesets | 51 |
| Viewing media list | 51 |
| | |
| Restoring file-level backups | 52 |
| About restoring file-level backups | 52 |
| Restoring file-level backups using shared network drives | 52 |
| Setting up a network share | 52 |
| Restoring data | 53 |
| Restoring file-level backups using the native Plug-in for FileSystem | 55 |
| Prerequisites | 55 |
| Restoring data | 55 |
| Using additional features available on the Choose Saveset page | 56 |
| Filtering the saveset list | 56 |
| Searching for files in savesets | 57 |
| Viewing media list | 57 |
| | |
| Troubleshooting | 58 |
| Common errors | 58 |
| SOAP messages | 61 |
| VDDK logs | 62 |
| | |
| About Dell | 63 |
| Contacting Dell | 63 |
| Technical support resources | 63 |

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 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® ESX® 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.

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, while 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 ESX 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 ESX 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 ESX Server or protect multiple ESX 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 VMware® ESX® and VMware® ESXi™ Servers as well as 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 can 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 confidently 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 to give you more efficient backups and reduced network demands.

The Plug-in for VMware also preserves the use of VMware vSphere® Thin Provisioning feature during virtual machine recovery operations to control storage costs and protect data 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
- Reduces backup windows by backing up multiple virtual machines in parallel
- Reduces network and storage requirements for NTFS basic disks 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 ESX 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 vCenter and ESX 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/>.

IMPORTANT: Starting with 10.0, NetVault Backup provides a web-based user interface to configure, manage, and monitor your NetVault Backup system and installed plug-ins. The procedures described in this guide are intended for the new NetVault WebUI. For procedures based on the NetVault Backup Console (user interface available with NetVault Backup 9.x), see the documentation for an earlier version of the plug-in.

- **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](#)
- [Deploying the plug-in](#)
- [Prerequisites](#)
- [Installing the plug-in](#)
- [Removing the plug-in](#)

About plug-in editions


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

- **ESX Server Edition:** The ESX Server Edition allows you to add one or more standalone ESX 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 ESX 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 ESX or 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 ESX or vCenter Servers to a client that runs the Enterprise Edition of the plug-in.

Deploying the plug-in

You can deploy the Plug-in for VMware on any Windows- or Linux-based NetVault Backup Client that has access to the virtual disk storage. This type of deployment enables centralized backup administration and off loads backup I/O from the ESX Server. The ESX Server is only required to generate the snapshot checkpoint, which is recorded in the disk. This results in a minimal performance impact on the ESX Server. With a vCenter Server, the plug-in takes advantage of the Storage vMotion technology to track virtual machine migration and automatically back up a virtual machine from its current host.

Figure 1. Plug-in for VMware deployment



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 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*.

ⓘ **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.

- **Review virtual disk storage requirements:** Depending on your virtual setup, verify that the following requirements are met:
 - Verify that the virtual disks that you want to back up reside on VMFS datastores created on local storage, iSCSI, or FC SAN. The plug-in also supports virtual machines that use Raw Device Mapping (RDM) disks in virtual compatibility mode.
 - If the ESX/ESXi Server uses Fibre Channel SAN or iSCSI SAN to store the virtual machine disks, select a dedicated physical machine to host the plug-in.

With other types of storage devices (including the local storage on the ESX Server), a physical machine is not necessary. You can configure one of the virtual machines as a NetVault Backup Client to host the plug-in.

- With SAN storage, the NetVault Backup Client requires read and write access to the SAN LUNs that contain the VMFS volumes (datastores) with virtual disks or the virtual compatibility RDMs.
To ensure access, add the NetVault Backup Client to the same fabric zones to which the ESX Server belongs. Additionally, set the same host mode (connection type) on the NetVault Backup Client and the ESX Server. For example, if you have an IBM array with the host mode set to LNXCL, use the same setting 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.
- **Enable SSL certificate verification:** When communicating with the vCenter or ESX Server, the Plug-in for VMware can transfer data over an encrypted SSL connection or an unencrypted network channel.
 - You can use an unencrypted channel if there is a trusted environment with fully isolated transmissions between the ESX Server and NetVault Backup Client. By disabling SSL, you can avoid the performance overhead associated with encryption.
 - To secure the communication channel, enable SSL certificate verification in your virtual environment. This option is disabled by default.
- **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 the 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 the advanced transport modes may fail. If the **libexpat** library is installed on a client, the plug-in automatically detects it, and creates the symbolic link "**libexpat.so.0**" to the library during installation.

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, create the symbolic link manually. If binary compatibility is present, you can also use a higher version of the library.

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 **/lib**, type:


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

To install the plug-in

- 1 In the Navigation pane, click **Guided Configuration**, and then on the **NetVault Configuration Wizard** page, click **Install Plugins**. In the **NetVault Backup Clients** list, select the clients on which you want to install the plug-in.

— or —

In the Navigation pane, click **Manage Clients**. In the **NetVault Backup Clients** list, select the client on which you want to install the plug-in, and click **Manage**. At the lower-right corner of the **Installed Software** table, click the **Install Plugin** button ().


 **NOTE:** You can use the configuration wizard to install the plug-in on multiple clients at the same time (if the selected clients are all of the same type). When you select multiple clients, verify that the plug-in binary file is compatible with the OS and platforms of the target clients.

From the **Manage Clients** page, you can only select one client for installing the plug-in.

- 2 Click **Choose Plug-in File**, and in the browse window, navigate to the location of the “.npk” installation file for the plug-in (on the installation CD or the directory to which the file was downloaded from the website).
- 3 Select the file named **vmw-w.x.y.z-<platform>.npk** (where **w.x** represent the version number, **y** represents the patch level, and **z** represents the build number), and click **Next** to begin installation.
- 4 After the plug-in is installed successfully, a message is displayed.

Removing the plug-in

To remove the plug-in

- 1 In the Navigation pane, click **Manage Clients**.
- 2 In the **NetVault 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 confirmation dialog box, click **Remove**.

Configuring the plug-in

- [Adding a server](#)
- [Configuring default settings](#)
- [Configuring CBT](#)
- [Enabling virtual machine locking](#)
- [Understanding backup consistency states](#)

Adding a server

Before you can start using the plug-in for backups and restores, you must add the applicable ESX/ESXi and vCenter Servers to the plug-in, and configure a user account to log on to the server.

- **Standalone ESX/ESXi Server setup:** In a VMware environment consisting of one or more ESX or 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. Select **VMware Plugin**, and in the **Actions** list, click **Add Server**.
- 4 Configure the following settings.

Table 1. Add Server

| Option | Description |
|-------------|--|
| Server Name | In a standalone server setup, type a name for the ESX or ESXi Server. In a cluster setup, type a name for the vCenter Server. The server name must be unique. If you do not provide a name, the plug-in uses the Server Address as the server name. NOTE: You cannot change the server name once it is configured. |

Table 1. Add Server

| Option | Description |
|----------------|--|
| | NOTE: When adding an ESX/ESXi 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 because the restore path is different. |
| Server Address | In a standalone server setup, type the Fully Qualified Domain Name (FQDN) of the ESX or ESXi Server. In a cluster setup, type the FQDN of the vCenter Server. You can also specify the IP address of the server, but using the FQDN is the preferred method. |
| User | 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. |
| Password | Type the password for the user account. |

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

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

Reconfiguring a server

To reconfigure an existing 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 ESX/ESXi or vCenter Server, and in the **Actions** list, click **Edit Server**.
- 4 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.
- 5 Click **OK** to save the new settings.

Removing a server

To remove 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 ESX/ESXi or vCenter Server, and in the **Actions** list, click **Remove Server**.
- 4 In the confirmation dialog box, click **OK**.

Configuring default settings

To configure the default settings for Plug-in for VMware

- 1 Start the Backup Job Wizard, and click **Create New** next to the **Selections** list. Open the NetVault Backup Server. Select **VMware Plugin**, and in the **Actions** list, click **Configure**.

— or —

In the Navigation pane, click **Change Settings**. On the **Configuration** page, click **Server Settings**, and then click **Plugin Options**.

- 2 Configure the following settings.

Table 2. Default settings for Plug-in for VMware

| Option | Description |
|----------------|--|
| Transport Mode | <p>The Plug-in for VMware supports the following methods for accessing virtual machine disks:</p> <p>SAN mode: The SAN transport mode supports virtual machine disks stored on Fibre Channel SAN, iSCSI SAN, or Serial Attached SCSI (SAS) storage devices. This transport mode requires a physical NetVault Backup Client. The SAN mode allows you to completely offload backups to a physical NetVault Backup Client.</p> <p>To use this transport mode, select san. Verify that the NetVault Backup Client on which the plug-in is installed has SAN connectivity to the VMFS LUN. On new plug-in installations, the san transport mode is selected by default.</p> <p><i>Best practices when using SAN transport for backup and restore</i></p> <ul style="list-style-type: none">• For array-based storage, use the SAN transport mode when the plug-in is running on a physical machine.• During restores, the SAN transport mode offers the best performance on thick disks. On thin disks, NBD and NBDSSL transport modes are faster than the SAN mode.• During SAN restores, disable CBT on the virtual machine. <p>SCSI Hot-Add mode: The SCSI Hot-Add mode supports all types of storage devices. To use this transport mode, the NetVault Backup Client must be installed on a virtual machine. The SCSI Hot-Add mode does not require you to expose SAN LUNs to the NetVault Backup Client. This transport mode is not as efficient as the SAN mode, which does not cause any overhead on the ESX Server host, but it is more efficient than the LAN mode.</p> <p>To use this transport mode, select hotadd.</p> <p><i>Requirements for SCSI Hot-Add mode</i></p> <p>To use the SCSI Hot-Add mode, the following requirements must be met:</p> <ul style="list-style-type: none">• The NetVault Backup Client must be installed on a virtual machine (that is, a backup proxy).• The backup proxy virtual machine must reside on the same datacenter as the target virtual machines (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. <p>The hotadd mode supports only LSI SCSI controller; it does not support Paravirtual SCSI controllers.</p> |

Table 2. Default settings for Plug-in for VMware

| Option | Description |
|-------------------------------------|---|
| Fallback Transport Mode | <p>LAN mode: The LAN transport mode supports all types of storage devices. You can use the LAN transport mode if the ESX/ESXi Server uses local storage devices or NAS to store its virtual machine disks. This transport mode does not require a physical NetVault Backup Client; you can use this transport mode if the NetVault Backup Client is installed on a virtual machine.</p> <p>The LAN mode uses the Network Block Device (NBD) protocol to access virtual disks. The ESX Server host reads data from the storage device, and sends the read data across a network channel to the plug-in.</p> <p>To use the LAN transport mode, select nbid or nbidssl:</p> <ul style="list-style-type: none"> • The nbid transport mode performs unencrypted data transfers and can be used if the ESX/ESXi Server and Plug-in for VMware reside on a secure isolated network. nbid is faster than nbidssl, and takes fewer resources of the ESX Server and NetVault Backup Client. • The nbidssl transport mode performs encrypted data transfers and can be used to protect sensitive data. <p>Auto mode: To automatically select the most suitable transport mode, select Auto.</p> <p>If your setup includes a combination of SAN, NAS, or locally attached SCSI storage for virtual machines, you can configure a fallback transportation mode that is used when the primary transportation mode fails.</p> <p>Depending on the type of storage your ESX or ESXi Server has access to, select san, hotadd, nbid, or nbidssl. If no alternative is available, select none. On new plug-in installations, the fallback transport mode is set to nbid by default.</p> |
| Default Inventory View | <p>The 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. <p>This 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 ESX Server setup, the virtual machines are displayed under the individual hosts.</p> <p>The Hosts and Clusters view does not display the folders created on the vCenter Server.</p> <ul style="list-style-type: none"> • 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 ESX/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 in the Job Monitor window. 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> |

Table 2. Default settings for Plug-in for VMware

| Option | Description |
|-------------------|--|
| Read Block Size | 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. |
| 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 the 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 <NetVault Backup home>\tmp on Windows and <NetVault Backup home>/tmp 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 Client.</p> |

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

Configuring CBT

The Changed Block Tracking (CBT) feature allows virtual machines to keep track of changed disk sectors. When CBT is enabled on a virtual machine, a “Change ID” is assigned to each disk during the snapshot generation activity. 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 RDM virtual disks, Virtual compatibility RDM (Independent Disks), or Virtual disks attached to a shared virtual SCSI bus.

Enabling CBT on a virtual machine

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. This section describes the procedure for enabling CBT on individual virtual machines.

- ① **NOTE:** You can also use the backup option “Enable Changed Block Tracking for VM” to automatically enable tracking on all virtual machines that are included in a backup job. For more information about this option, see [Other backup options for Plug-in for VMware](#).

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 ESX/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, and in the **Actions** list, click **Enable Change Block Tracking**.
This option is available to only those virtual machines on which CBT is disabled.
- 5 After the virtual machine is reconfigured, a message is displayed. Click **OK** to close the dialog box.

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.

To disable 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 ESX/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, and in the **Actions** list, click **Disable Change Block Tracking**.
This option is available to only those virtual machines on which CBT is enabled.
- 5 After the virtual machine is reconfigured, a message is displayed. Click **OK** to close the dialog box.

Manually disabling CBT on a virtual machine

The plug-in versions earlier than 2.7 do not provide an option to disable CBT from the NetVault Backup Console. On such installations, if you want stop using CBT, you must manually disable this feature.

To manually disable CBT on a virtual machine

- 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 virtual machine locking

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 or it can create orphaned virtual disks on the datastores. To temporarily disable migration, the virtual machines can be locked before a job starts and unlocked when the job is completed.

By default, the locking mechanism is disabled and it is not used during backups and restores.

To enable virtual machine locking

- 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 In the `[Custom:GetMigrationLock]` section, set the `Value` option to `True`:

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

- 3 Save the file.

Understanding backup consistency states

The Plug-in for VMware supports the following snapshot consistency states.

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.

On Linux-based Guest OS, the Plug-in for VMware only supports crash-consistent snapshots.

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.

This consistency state is only supported on Windows-based Guest OS.

Requirements for file-system-consistent snapshots

Table 3. Requirements for file-system-consistent snapshots

| Guest OS version | ESX Server version | Requirements |
|-------------------------------|-----------------------------|--|
| Any supported Windows version | ESX Server 3.5 U2 and later | Install VSS components on the virtual machine as part of updated VMware Tools. |
| Any supported Windows version | ESX Server prior to 3.5 U2 | Install the SYNC driver included in the VMware Tools on the virtual machine. |

NOTE: The VMware SYNC driver can also be installed 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.

NOTE: Creating quiesced snapshots using SYNC driver requires waiting for I/O to drain in the Guest OS. This action can affect time-sensitive applications.

VSS-based application-consistent snapshots

When backing up a Windows-based virtual machine, the Plug-in for VMware uses VSS to perform application-level quiescing before generating a snapshot. 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.

This consistency state is only supported on Windows 2008 and Windows 2003.

Requirements for VSS-based application-consistent snapshots

Table 4. Requirements for VSS-based application-consistent snapshots

| Guest OS version | ESX Server/vSphere version | Requirements |
|---------------------------------|-----------------------------|--|
| Windows 2008 Windows 2008 R2 | vSphere 4.1 and later | <ul style="list-style-type: none">• Install VSS components on the virtual machine as part of updated VMware Tools.• Enable the UUID attribute for the virtual machine, if it is not enabled. To enable this attribute, see Enabling UUID on Windows 2008 Guest OS. (This attribute is enabled by default for virtual machines created on ESX 4.1 and later.)• Ensure that the virtual machines only use SCSI disks and have enough free SCSI slots as the number of disks. Application-consistent snapshots are not supported on virtual machines that use IDE disks.• Ensure that the virtual machines do not use dynamic disks. |
| Windows 2003 | ESX Server 3.5 U2 and later | <ul style="list-style-type: none">• Install VSS components on the virtual machine as part of updated VMware Tools. |

Enabling UUID on Windows 2008 Guest OS

To enable UUID on Windows 2008 Guest OS

- 1 Start the vSphere Client.
- 2 Log on to the vCenter Server, and select **VMs and Templates**.
- 3 Click the **Virtual Machines** tab.
- 4 Power off the virtual machine (right-click the virtual machine, and select **Power > Power Off**).
- 5 Right-click the virtual machine, and select **Edit Settings**.
- 6 Click the **Options** tab. In the settings column, select **General**, and then click **Configuration Parameters**.
- 7 In the dialog box that appears, click **Add Row**, and configure the following settings:
 - **Name:** Type `disk.EnableUUID`.
 - **Value:** Type `TRUE`.Click **OK**.
- 8 Save the settings and power on the virtual machine.

Defining a backup strategy

- [About backup methods and types](#)
- [Defining a backup and recovery strategy](#)

About backup methods and types

The Plug-in for VMware supports the following 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.

This backup method is available to both Linux- and Windows-based virtual machines.

Image-level backups with CBT

The Plug-in for VMware supports the following types of CBT-based image-level backups.

- **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.

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 1. 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 | 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. |

Image-level backups without CBT

The Plug-in for VMware only supports Full image-level backups of virtual machines that do not use CBT.

- 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.

Important notes

The following table describes how the plug-in handles backups and restores of RDM disks.

Table 2. 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. |

Table 2. Considerations for RDM disks

| RDM compatibility mode | VADP-based backups and restores |
|----------------------------|--|
| 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><i>To omit the RDM disks in virtual compatibility mode</i></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. |

File-level backups

This 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.

Defining a 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](#)
- [Resetting CBT on a virtual machine](#)
- [Monitoring job progress](#)

Performing image-level backups

To perform an image-level backup

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

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

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

- a Click **Create New** to open the **NetVault Backup Selections** page.
- b Open the NetVault Backup Client on which the plug-in is installed, and then open **VMware Plugin**.
- c Open the applicable ESX/ESXi or vCenter Server. Depending on your VMware setup and inventory view type, open the applicable container nodes.

NOTE: To switch between the two views, select the ESX/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 within a resource pool, select the corresponding resource pool node. Similarly, to back up all virtual machines hosted on an ESX or 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, Resource Pool, and other nodes), 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.
Alternatively, 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. Note the following:
 - 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.

Alternatively, select the container node (for example, Datacenter, Resource Pool, ESX 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](#).

NOTE: 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. Click **Save** to close the dialog box.

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

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

- a Click **Create New** to open the **VMware Plugin Backup Options** page.

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.

- b Under **Backup Type**, select one of the following options.

Table 1. 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](#).

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 following error occurs:

- **vSphere 4.0:** On vSphere 4.0, Incremental Backups include only those disk sectors that have changed since the last Incremental Backup instead of including all sectors that have changed since the virtual machine was reverted. This omission can result in data corruption or data loss without your knowledge as vSphere does not report this error. The jobs can succeed or fail, and in some cases succeed after multiple retries. Therefore, it is necessary to avoid this situation.
- **vSphere 4.1 and later:** On vSphere 4.1 and later, 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.

ⓘ IMPORTANT: 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, ESX 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 2. 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 3. Other backup options for Plug-in for VMware

| Option | Description |
|--------------------------------------|--|
| Enable Changed Block Tracking for VM | 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. To enable CBT on all virtual machines included in the backup job, select this check box. Before you select this check box, verify that the following requirements are met: <ul style="list-style-type: none"> • The virtual machines are managed by ESX/ESXi 4.0 or later. • The virtual machines use virtual hardware version 7 or later. • The virtual machines do not use Physical compatibility RDM virtual disks, Virtual compatibility RDM (Independent Disks), or Virtual disks attached to a shared virtual SCSI bus. If you select this check box and the plug-in fails to change this setting on a virtual machine, a warning message is logged. 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. NOTE: While this setting is applied to all virtual machines included in a backup job, the plug-in also provides you an option to enable CBT on individual virtual machines. For more information about enabling CBT on specific virtual machines, see Enabling CBT on a virtual machine . |

Table 3. 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 ESX 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, which reduce 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. Currently, ABM is only supported on the NTFS file system on basic disks.</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>You can enable ABM for jobs that include both NTFS disks and non-NTFS disks, but it is only applied to the NTFS disks.</p> |
| 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> |

Table 3. Other backup options for Plug-in for VMware

| Option | Description |
|--------|--|
| | <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 machine that are managed by the same ESX 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 additional 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> <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> |

- e Click **Save**, and in the **Create New Set** dialog box, type a name for the set.
- 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**. You can monitor the job progress from the **Job Status** page and view the logs from the **View Logs** page.

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

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

Additional notes

- When you select a container node (for example, Datacenter, Resource Pool, or ESX/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 ESX 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.

- If Fault Tolerance (FT) is enabled for a virtual machine, the plug-in does not allow you to select it for a backup.

If this feature is enabled after including a virtual machine in a backup job, that virtual machine is not backed up and the message "Cannot backup VM – Fault Tolerance is enabled for the VM" is added to the job logs.

- 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.

- 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 as soon as the backup 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, including power-on, resume after suspend, migrate, or snapshot create, delete, or revert operations.
- During an Incremental or Differential Backup, the plug-in does not attempt to enable CBT for any virtual machine included in the backup, unless you have selected the **Enable Change Block Tracking for VM** check box.

- If you add a new virtual machine to the Backup Selection Set (or to an ESX 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.
- 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 the **Enable Changed Block Tracking for VM** check box for each job.

To stop using CBT, you must disable this feature. For instructions, see [Disabling CBT on a virtual machine](#) or [Manually disabling CBT on a virtual machine](#).

- 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 4. Backup selection tree icons





























| Icon | Description |
|---|--|
|  | vCenter Server |
|  | Datacenter Server |
|  | ESX or ESXi Server Cluster |
|  | Closed folder |
|  | Open folder |
|  | ESX or ESXi Server |
|  | ESX or ESXi Server in maintenance mode |
|  | Inaccessible ESX or 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 |

Table 4. Backup selection tree icons

| Icon | Description |
|---|---|
|  | 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 |

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.

To reset 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 ESX/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, 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, we recommend 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 backup.
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

To monitor the progress of a job

- 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](#)
- [Additional notes](#)

Performing file-level backups

To perform a file-level backup

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

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

- 3 In the **Selections** list, select an existing Backup Selection Set, or complete the following steps to create a set.
 - a Click **Create New** to open the **NetVault Backup Selections** page.
 - b Open the NetVault Backup Client on which the plug-in is installed, and then open **VMware Plugin**.
 - c Open the applicable ESX/ESXi or 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.

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. 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.
 - ❗ **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.
 - f Click **Save**, and in the **Create New Set** dialog box, type a name for the set. Click **Save** to close the dialog box.
 - ❗ **NOTE:** A set name can contain alphanumeric and non-alphanumeric characters, but it cannot include non-English characters. On Linux OS, the names can have a maximum of 200 characters. On Windows OS, there is no length restriction. However, a maximum of 40 characters is recommended on all platforms.
- 4 In the **Plugin Options** list, select an existing Backup Options Set, or complete the following steps to create a set.
- a Click **Create New** to open the **File System Plugin Backup Options** page.
 - ❗ **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.
 - b Configure the following settings.

Table 1. 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.
- 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**. You can monitor the job progress from the **Job Status** page and view the logs from the **View Logs** page.

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

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

Backup selection tree icons

Table 2. Backup selections tree icons





























| Icon | Description |
|---|--|
|  | vCenter Server |
|  | Datacenter Server |
|  | ESX or ESXi Server Cluster |
|  | Closed folder |
|  | Open folder |
|  | ESX or ESXi Server |
|  | ESX or ESXi Server in maintenance mode |
|  | Inaccessible ESX or 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) |

Table 2. Backup selections tree icons

| Icon | Description |
|---|--|
|  | 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 |

Additional notes

Manually removing the snapshot and mount folder

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” for 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, 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 manually remove the snapshot and mount folder

- 1 If `nvvmwmount.exe` is running, unmount the virtual machines that are currently mounted for browse or backup operations. Stop the NetVault Backup Service. This step is only applicable to 2.1.55 and earlier versions of the Plug-in for VMware.
- 2 If the Working Directory contains the mount folder for the virtual machine, remove it.

- 3 If you were using an advanced transport mode, such as **san** or **hotadd**, navigate to the <system_drive>/windows/temp/vmware-system directory. 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 ESX 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”) may be displayed 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.

- 6 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.
- 7 If some unwanted redo logs are still present on the datastore, perform step 6 again. 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](#)
- [Using other restore procedures](#)
- [Performing file-level restores from image-level backups](#)
- [Restoring virtual machine disk and configuration files](#)
- [Using additional features available on the Choose Saveset page](#)

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 ESX 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 allows you to 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 step 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:

- If the virtual machine that you want to restore still exists on the host server, delete it.
- 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

- 1 In the Navigation pane, click **Create Restore Job**.
- 2 In the saveset table, select the saveset that you want to use, and click **Next**.

The table displays the saveset name (Job Title and Saveset ID), creation date and time, and saveset size. Note the following:

- The list is sorted alphabetically by saveset name. You can sort the list by a different column or reverse the sort order by clicking the column heading. The arrowhead next to the column name indicates the sort order.
- You can use one or more filters to display specific savesets on this page. You can also search for a data item in savesets and view the media list for a saveset. For more information about the additional features, see [Using additional features available on the Choose Saveset page](#).
- When you select a saveset, the following details are displayed in the **Saveset Information** area: Job ID, Job Title, name of the NetVault Backup Server, name of the client from which the data was backed up, plug-in used to create the saveset, saveset creation date and time, saveset retirement setting, whether Incremental Backup or not, whether Archive or not, and saveset size.

- 3 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.

- 4 Click **Edit Plugin Options**, and in the **Plug-in for VMware Restore Options** dialog box, configure the following settings.

Table 1. Options for full virtual machine recovery

| Option | Description |
|-------------------------------------|--|
| Recover VM | To restore a full virtual machine or one or more virtual disks, select this option. |
| Relocate to 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 ESX 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>The restore fail if the datastore is invalid or inaccessible or does not have sufficient space to store virtual machine files.</p> |
| Allocate to Different Resource Pool | <p>A resource pool represents the CPU and memory resources that are available to an ESX 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>/Pool-A/Pool-B/.../Pool<n></pre> <p>Here, Pool<n> is the target resource pool, Pool-A is the child of the root resource pool, Pool-B is the child of Pool-A, 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. |

Table 1. Options for full virtual machine recovery

| Option | Description |
|---|--|
| Perform COS-Less Restore <i>(Legacy backups only)</i> | For savesets created with Plug-in for VMware 1.x, the Perform COS-Less Restore check box must be selected to specify an alternate datastore or resource pool. If you do not select this check box, the plug-in ignores the settings for the Relocate to Alternate Datastore and Allocate to Different Resource Pool options for such savesets. NOTE: This additional check box is only available to legacy backups created with Plug-in for VMware 1.x. The plug-in no longer supports the COS-based and COS-less restore types. The later versions of the plug-in ignore this option during a restore. |

- 5 Click **Save** to save the settings, and then click **Next**.
- 6 In **Job Name**, specify a name for the job.

Assign a descriptive name that allows you to easily identify the job for monitoring its progress. The job name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. There is no length restriction. However, a maximum of 40 characters is recommended on all platforms
- 7 In the **Target Client** list, the client from which data was backed up is selected by default. Do not change this setting.
- 8 Select or create the **Schedule Set**, **Restore Source Set**, and **Advanced Options Set**. For more information about these sets, see the *Dell NetVault Backup Administrator's Guide*.
- 9 Click **Submit** to submit the job for scheduling.

You can monitor the job progress from the **Job Status** page and view the logs from the **View Logs** page. For more information about these functions, see the *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.
- If you select one or more virtual disks instead of the virtual machine, the plug-in overwrites the contents of the existing ".vmdk" file. If the ".vmdk" file is not available on the vCenter Server, the plug-in reports an error.
- 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 2. Restore selection tree icons

| Icon | Description |
|---|------------------------------|
|  | Virtual machine |
|  | 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. In most cases, you can ignore this message and select the normal boot option.

Using other restore procedures

Relocating a virtual machine to an alternate ESX Server host

To relocate a virtual machine to an alternate ESX Server host

- 1 Complete steps 1 through 4 in the section [Restoring data](#).
- 2 Under Alternate Target ESX Server Details, configure the following settings:

Table 3. Relocate a virtual machine to alternate ESX Server

| Option | Description |
|-----------|---|
| Host | Type the IP address or DNS name of the alternate ESX Server host. |
| User Name | Specify a user account that can be used to log on to the ESX Server. The user account must have privileges to register or create virtual machines |
| Password | Type the password for the user account. |

- 3 Complete steps 5 through 9 in the section [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 right-hand panel, 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 rename a virtual machine during restore

- 1 Complete steps 1 through 3 in the section [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.

- 3 In the **Restore Rename** dialog box, provide the following information:
 - **Rename To:** Select this 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.

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

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.

- 4 Complete steps 4 through 9 in the section [Restoring data](#).

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

- 5 When you start the renamed virtual machine, 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

To perform file-level restores from image-level backups

- 1 In the Navigation pane, click **Create Restore Job**.
- 2 In the saveset table, select the saveset that you want to use, and click **Next**.
The table displays the saveset name (Job Title and Saveset ID), creation date and time, and saveset size. Note the following:
 - The list is sorted alphabetically by saveset name. You can sort the list by a different column or reverse the sort order by clicking the column heading. The arrowhead next to the column name indicates the sort order.
 - You can use one or more filters to display specific savesets on this page. You can also search for a data item in savesets and view the media list for a saveset. For more information about the additional features, see [Using additional features available on the Choose Saveset page](#).
 - When you select a saveset, the following details are displayed in the **Saveset Information** area: Job ID, Job Title, name of the NetVault Backup Server, name of the client from which the data was backed up, plug-in used to create the saveset, saveset creation date and time, saveset retirement setting, whether Incremental Backup or not, whether Archive or not, and saveset size.
- 3 On the **Create Selection Set** page, open the applicable container nodes (that is, the vCenter, Datacenter, ESX Host Cluster, and ESX Server) to list the virtual machine images stored in the saveset.
- 4 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.
- 5 Select the files and directories that you want to restore.
- 6 Click **Edit Plugin Options**, and in the **Plug-in for VMware Restore Options** dialog box, configure the following settings.

Table 4. Options for file-level restores from image-level backups

| Option | Description |
|---------------------------|---|
| Restore Files from Backup | 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. |

- 7 Click **Save** to save the settings, and then click **Next**.

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










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

- 9 In the **Target Client** list, the client from which data was backed up 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 Click **Submit** to submit the job for scheduling.

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

Restore selection tree icons

Table 5. Restore selection tree icons

| Icon | Description |
|---|------------------------------|
|  | Virtual machine |
|  | 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

- 1 In the Navigation pane, click **Create Restore Job**.
- 2 In the saveset table, select the saveset that you want to use, and click **Next**.

The table displays the saveset name (Job Title and Saveset ID), creation date and time, and saveset size. Note the following:

- The list is sorted alphabetically by saveset name. You can sort the list by a different column or reverse the sort order by clicking the column heading. The arrowhead next to the column name indicates the sort order.
 - You can use one or more filters to display specific savesets on this page. You can also search for a data item in savesets and view the media list for a saveset. For more information about the additional features, see [Using additional features available on the Choose Saveset page](#).
 - When you select a saveset, the following details are displayed in the **Saveset Information** area: Job ID, Job Title, name of the NetVault Backup Server, name of the client from which the data was backed up, plug-in used to create the saveset, saveset creation date and time, saveset retirement setting, whether Incremental Backup or not, whether Archive or not, and saveset size.
- 3 On the **Create Selection Set** page, select the image from which you want to restore the disk and configuration files.
 - 4 Click **Edit Plugin Options**, and in the **Plug-in for VMware Restore Options** dialog box, configure the following settings.

Table 6. Options for restoring virtual machine disk and configuration files from image-level backups

| Option | Description |
|---------------------------|---|
| Restore Files from Backup | 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. |

Table 6. Options for restoring virtual machine disk and configuration files from image-level backups

| Option | Description |
|------------------------------|---|
| Split Restored Virtual Disks | <p>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.</p> <p>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 <code>vmware.cfg</code> file:</p> <ol style="list-style-type: none"> 1 Open the <code>vmware.cfg</code> file in any 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 Add the following lines: <pre>[Custom:RestoreSplitVMDK] Value=TRUE</pre> 3 Save the file. <p>The setting for this option in the Plug-in for VMware Restore Options dialog box overrides the setting in the <code>vmware.cfg</code> file.</p> |

- 5 Click **Save** to save the settings, and then click **Next**.
- 6 In **Job Name**, specify a name for the job.

Assign a descriptive name that allows you to easily identify the job for monitoring its progress. The job name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. There is no length restriction. However, a maximum of 40 characters is recommended on all platforms
- 7 In the **Target Client** list, the client from which data was backed up is selected by default. Do not change this setting.
- 8 Select or create the **Schedule Set**, **Restore Source Set**, and **Advanced Options Set**. For more information about these sets, see the *Dell NetVault Backup Administrator's Guide*.
- 9 Click **Submit** to submit the job for scheduling.
- 10 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 7. Restore selection tree icons







| Icon | Description |
|---|------------------------------|
|  | Virtual machine |
|  | Configuration file |
|  | NVRAM file |
|  | Virtual machine disk file |
|  | Disk extent information file |

Table 7. Restore selection tree icons

| Icon | Description |
|---|--------------|
|  | 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.” Additionally, 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 disk(s).

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. Follow the instructions to complete the virtual machine recovery. When the “Select source type” prompt appears, select “VMware Workstation or other VMware virtual machine.” Additionally, 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.

Using additional features available on the Choose Saveset page

This section describes how to use the additional features available on the **Create Restore Job – Choose Saveset** page. These features are common to all NetVault Backup plug-ins.

Filtering the saveset list

By default, the saveset table on the **Create Restore Job – Choose Saveset** page lists all available savesets. You can use the following filters to display specific savesets on this page.

Table 8. Saveset filters

| Filter | Description |
|-------------|---|
| Client | <p>Displays savesets created for particular clients.</p> <p><i>To use this filter</i></p> <ul style="list-style-type: none"> • Click the Client box, and in the Choose Client dialog box, select the applicable clients. Click OK to close the dialog box. <p>The default selection is Any.</p> |
| Plugin Type | <p>Displays savesets created using a particular plug-in.</p> <p><i>To use this filter</i></p> <ul style="list-style-type: none"> • Click the Plugin Type box, and in the list, select the applicable plug-in. <p>The default selection is Any.</p> |
| Date | <p>Displays savesets created during a specified period.</p> <p><i>To use this filter</i></p> <ul style="list-style-type: none"> • Click the Date box, and 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. <p>The default selection is Any.</p> |
| Job | <p>Displays savesets created for particular job IDs.</p> <p><i>To use this filter</i></p> <ul style="list-style-type: none"> • Click the Job box, and in the Choose Job dialog box, select the applicable jobs. Click OK to close the dialog box. <p>The default selection is Any.</p> |

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 a saveset or browsing through its 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. Select the items you want to restore. You can only restore items from one saveset.

Viewing media list

To view media list for a saveset

- 1 On the **Create Restore Job – Choose Saveset** page, select the applicable saveset.
- 2 In the **Saveset Information** area, click **Media List**.
- 3 In the dialog box that appears, you can view the data and index segment details. For each data segment, you can view the media label, media group, offset, segment size, and media location. For index segments, you can view the media label, media group, and media location.
- 4 Click **Close** to close the dialog box.

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](#)
- [Using additional features available on the Choose Saveset page](#)

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

- 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.

i | **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 Click **OK** to save the connection details and close the dialog box.

Restoring data

To restore file-level backups

- 1 In the Navigation pane, click **Create Restore Job**.
- 2 In the saveset table, select the saveset that you want to use, and click **Next**.

The table displays the saveset name (Job Title and Saveset ID), creation date and time, and saveset size. Note the following:

- The list is sorted alphabetically by saveset name. You can sort the list by a different column or reverse the sort order by clicking the column heading. The arrowhead next to the column name indicates the sort order.
 - You can use one or more filters to display specific savesets on this page. You can also search for a data item in savesets and view the media list for a saveset. For more information about the additional features, see [Using additional features available on the Choose Saveset page](#).
 - When you select a saveset, the following details are displayed in the **Saveset Information** area: Job ID, Job Title, name of the NetVault Backup Server, name of the client from which the data was backed up, plug-in used to create the saveset, saveset creation date and time, saveset retirement setting, whether Incremental Backup or not, whether Archive or not, and saveset size.
- 3 On the **Create Selection Set** page, select the files and directories that you want to restore.

i | **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), but for a Dump Type Incremental or Differential Backup, the plug-in only displays those files that were backed up in the selected saveset.
 - 4 Select the file or directory that you want to restore, and in the **Actions** list, click **Rename**.
 - 5 In the **Relocate** box, type the network share path.

- i | **NOTE:** If you do not perform steps 4 and 5, 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 1. Plug-in for FileSystem restore options

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

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

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

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

- In the **Target Client** list, the client from which data was backed up 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*.
- Click **Submit** to submit the job for scheduling.

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

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

- 1 In the Navigation pane, click **Create Restore Job**.
- 2 In the saveset table, select the saveset that you want to use, and click **Next**.

The table displays the saveset name (Job Title and Saveset ID), creation date and time, and saveset size. Note the following:

- The list is sorted alphabetically by saveset name. You can sort the list by a different column or reverse the sort order by clicking the column heading. The arrowhead next to the column name indicates the sort order.
 - You can use one or more filters to display specific savesets on this page. You can also search for a data item in savesets and view the media list for a saveset. For more information about the additional features, see [Using additional features available on the Choose Saveset page](#).
 - When you select a saveset, the following details are displayed in the **Saveset Information** area: Job ID, Job Title, name of the NetVault Backup Server, name of the client from which the data was backed up, plug-in used to create the saveset, saveset creation date and time, saveset retirement setting, whether Incremental Backup or not, whether Archive or not, and saveset size.
- 3 On the **Create Selection Set** page, select the files and directories that you want to restore.
 - ① **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), but for a Dump Type Incremental or Differential Backup, the plug-in only displays those files that were backed up in the selected saveset.
 - 4 Click **Edit Plugin Options**, and configure the applicable settings. For more information, see [Plug-in for FileSystem restore options](#).
 - 5 Click **Save** to save the settings, and then click **Next**.
 - 6 In **Job Name**, specify a name for the job.

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

7 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 will be restored to the client on which the plug-in is installed.

8 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*.

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

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

Using additional features available on the Choose Saveset page

This section describes how to use the additional features available on the **Create Restore Job – Choose Saveset** page. These features are common to all NetVault Backup plug-ins.

Filtering the saveset list

By default, the saveset table on the **Create Restore Job – Choose Saveset** page lists all available savesets. You can use the following filters to display specific savesets on this page.

Table 2. Saveset filters

| Filter | Description |
|-------------|--|
| Client | Displays savesets created for particular clients. <i>To use this filter</i> <ul style="list-style-type: none">Click the Client box, and in the Choose Client dialog box, select the applicable clients. Click OK to close the dialog box. The default selection is Any. |
| Plugin Type | Displays savesets created using a particular plug-in. <i>To use this filter</i> <ul style="list-style-type: none">Click the Plugin Type box, and in the list, select the applicable plug-in. The default selection is Any. |
| Date | Displays savesets created during a specified period. <i>To use this filter</i> <ul style="list-style-type: none">Click the Date box, and 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. The default selection is Any. |
| Job | Displays savesets created for particular job IDs. <i>To use this filter</i> <ul style="list-style-type: none">Click the Job box, and in the Choose Job dialog box, select the applicable jobs. Click OK to close the dialog box. The default selection is Any. |

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 a saveset or browsing through its 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. Select the items you want to restore. You can only restore items from one saveset.

Viewing media list

To view media list for a saveset

- 1 On the **Create Restore Job – Choose Saveset** page, select the applicable saveset.
- 2 In the **Saveset Information** area, click **Media List**.
- 3 In the dialog box that appears, you can view the data and index segment details. For each data segment, you can view the media label, media group, offset, segment size, and media location. For index segments, you can view the media label, media group, and media location.
- 4 Click **Close** to close the dialog box.

Troubleshooting

- [Common errors](#)
- [SOAP messages](#)
- [VDDK logs](#)

Common errors

The following is a list of some common errors and their solution.

Table 1. Common errors and their solution

| Description | Symptom | Solution |
|---|---|---|
| The NetVault Backup Service fails to start on a Windows-based NetVault Backup Server. | Check the Windows Event Viewer to see if it displays the following message: PDT FATAL: lock file "postmaster.pid" already exists | 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. For more information, see https://support.software.dell.com/netvaultbackup/kb/122475 . |
| After restarting the machine, the NetVault Backup Service sometimes fails to start on a Windows-based NetVault Backup Server. | 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 | 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 from the Configurator. |
| The NetVault Backup Service starts, and then stops immediately on a Linux-based machine. | No error messages are displayed. | This issue can occur if the Postgres service cannot resolve the host name localhost , and fails to start. Check the <code>/etc/hosts</code> file, and if the file does not contain an entry for localhost , add the entry. |

Table 1. Common errors and their solution

| Description | Symptom | Solution |
|---|--|---|
| The Linux Hybrid version of the plug-in does not run on pure 64-bit operating system. | The plug-in fails to open and displays the message "No Error." | <p>Pure 64-bit operating systems do not install some common 32-bit libraries the plug-in requires.</p> <p>To correct this problem, do the following:</p> <ol style="list-style-type: none"> 1 Start a terminal session, and go to <code>/usr/netvault/bin</code>. 2 Run the following command to list the missing libraries: <pre>ldd nvvmware egrep -i "missing not found"</pre> 3 Install the missing 32-bit libraries using the applicable package manager. <p>NOTE: Some libraries are dynamically loaded from the path <code>\$NV_HOME/dynlib/vddk/</code>. If these libraries are listed as missing, no action is required.</p> |
| The plug-in crashes during a backup. | <p>During a backup, the plug-in crashes with the following error in the NetVault Backup logs:</p> <p>Backing up data of VM disk [XXXXXXX]xxx/xxxxxxxxx.vmdk' to media : "Job manager lost message channel unexpectedly."</p> <p>The trace logs include the following message:</p> <p>Cannot open library: libexpat.so.0: cannot open shared object file in vixDiskLib-16642.log.</p> | <p>To correct this problem, navigate to the <code>/usr/lib</code> directory, and create the following symbolic link:</p> <pre>ln -s /lib/libexpat.so.1.5.2 libexpat.so.0</pre> |
| CBT-based image-level backup fails. | <p>An image-level backup for which CBT has been enabled reports the following error and fails:</p> <p>Failed to get changed disk areas.</p> | <p>This issue occurs when a snapshot exists on a virtual machine when CBT is enabled. For snapshots that exist before enabling CBT, the <code>changed</code> parameter is not set. So, when the <code>QueryChangedDiskAreas</code> API is invoked, it returns an error. For more information, see http://kb.vmware.com/kb/1033816.</p> <p>To correct this problem, do either of the following:</p> <ul style="list-style-type: none"> • Ensure that no snapshots exist on the virtual machine before selecting the Enable Changed Block Tracking for VM check box. |

Table 1. Common errors and their solution

| Description | Symptom | Solution |
|--|---|--|
| 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. | <p>During a backup using the SAN transport mode, the following errors are reported:</p> <p>San transport error: I/O Operation failed.</p> <p>Error: One of the parameters supplied is invalid.</p> | <ul style="list-style-type: none"> Alternatively, use the Reset Change Block Tracking option if you are using Plug-in for VMware 2.7 or later. <p>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.</p> <p>These errors are caused by a fault in the Linux kernel when handling SCSI reservation conflicts.</p> <p>The errors are encountered on the following systems:</p> <ul style="list-style-type: none"> RHEL 6.1 and earlier versions SLES 11.1 and earlier versions <p>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.</p> |
| The file-level backup method does not back up data through a reparse point. | <p>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.</p> | <p>The VMware API does not support this type of backup.</p> <p>Workaround:</p> <ul style="list-style-type: none"> Install NetVault Backup Client on the Guest OS, and perform a backup using the Plug-in for FileSystem to include the reparse point directory. Alternatively, 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. | <p>A restore job fails with the following error:</p> <p>Log Message: Failed to restore VM disk data.</p> <p>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/VitrualDiskName.vmdk'</p> | <p>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 the restore completes successfully with this mode, it indicates SAN connectivity issues.</p> <p>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.</p> |

Table 1. Common errors and their solution

| Description | Symptom | Solution |
|--------------------|--|--|
| Restore job fails. | <p>A restore job fails with the following error:</p> <p>Log message: Failed to add VM to target server's inventory.</p> <p>Log context: Fault string is 'Invalid configuration for device '3'.</p> <p>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:</p> <pre><val xsi:type="LocalizedMethodFault"> <fault xsi:type="InvalidDeviceSpec"> <property>deviceChange[3]. device.backing.fileName </property> <deviceIndex>3</deviceIndex> </fault> <localizedMessage>Invalid configuration for device '3';.</localizedMessage> </val></pre> <p>If yes, it indicates that the backup contains an invalid configuration for a CD or DVD drive.</p> | <p>To correct this problem</p> <ol style="list-style-type: none"> 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: <pre>[Custom:ReconfigureVirtualCdromDevices] Value=True</pre> (The default value for this entry is False.) Save the file. 3 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. 4 After restoring the virtual machine, reset [Custom:ReconfigureVirtualCdromDevices] to False. |

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 ESX/ESXi Servers.

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

VDDK logs

The Plug-in for VMware uses the VMware VDDK API for image-level backups and restores and file-level browse and backup operations. To diagnose or troubleshoot errors encountered during such operations, Dell Software Technical Support may require the VDDK logs.

To generate VDDK logs

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

```
[Custom:GenerateVddkLogsOnDisk]
```

```
Value = True
```

By default, this value is set to **False**. It should only be changed under the direction of Dell Software Technical Support.

Save the file.

- 3 Retry the operation with the new setting. The plug-in generates VDDK log files entitled `vixDiskLib*.log` in the following directories:
 - **Windows:** `<SYSTEMDRIVE>\WINDOWS\Temp\vmware-SYSTEM`
 - **Linux:** `/tmp/vmware-root`
- 4 After capturing the logs, reset `[Custom:GenerateVddkLogsOnDisk]` to **False**.

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

Contacting Dell

Technical support:

[Online support](#)

Product questions and sales:

(800) 306-9329

Email:

info@software.dell.com

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 portal provides direct access to product support engineers through an online Service Request system.

The site enables you to:

- Create, update, and manage Service Requests (cases)
- View Knowledge Base articles
- Obtain product notifications
- Download software. For trial software, go to [Trial Downloads](#).
- View how-to videos
- Engage in community discussions
- Chat with a support engineer