NetBackup™ for VMware Administrator's Guide

Release 11.0

VERITAS
now part of Cohesity

Last updated: 2025-03-18

Legal Notice

Cohesity, Inc.
2625 Augustine Drive
Santa Clara, CA 95054

http://www.veritas.com

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

https://www.veritas.com/support

You can manage your Cohesity account information at the following URL:

https://my.veritas.com

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Cohesity website:

https://sort.veritas.com/documents

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

NB.docs@veritas.com

You can also see documentation information or ask a question on the Cohesity community site:

http://www.veritas.com/community/

Cohesity Services and Operations Readiness Tools (SORT)

Cohesity Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1

Introduction

This chapter includes the following topics:

About NetBackup for VMware

NetBackup for VMware provides backup and restore of the VMware virtual machines that run on VMware ESX servers. NetBackup for VMware takes advantage of VMware vStorage APIs for data protection. The backup process is off-loaded from the ESX server to a VMware backup host.

NetBackup for VMware does the following:

- Performs off-host backup of virtual machines (NetBackup client software is not required on the virtual machine). Off-host backup reduces the backup processing load on the VMware host.
- Increases the backup speed as compared to standard file-order backup methods, if the virtual machine is heavily populated with small files.

- Automatically creates quiesced snapshots using VSS (Windows only). Creates quiesced snapshots on Linux if snapshot quiesce is enabled in the Linux guest OS.
- Uses snapshot technology to keep virtual machines 100% available to users.
- Supports VMware vSphere and vCloud Director.
- Performs full backups and incremental backups, including block-level incrementals.
- Backs up the full virtual machine.
- Backs up the virtual machines even when they are turned off.
- Can restore selected files from the backup.

About the virtual machine backups that include database data

When NetBackup backs up a VMware virtual machine, database data in the virtual machine is backed up with the rest of the virtual machine. NetBackup allows the recovery of individual database files from the backup. This feature supports Microsoft Exchange Server, SQL Server, and SharePoint Server.

Note the following:

- To enable restore of individual database files, a NetBackup Windows client must be installed in the virtual machine during the backup.
- NetBackup uses Windows Volume Shadow Copy Service (VSS) to quiesce the database before it creates a snapshot of the virtual machine.
- A full backup is performed of the database data with each backup job, even if the policy schedule is incremental.

See “Application protection options (VMware)” on page 96.

About the NetBackup appliance as a VMware backup host

The NetBackup appliance uses the VMware policy type to back up VMware virtual machines.

The following topics contain notes on the appliance as the backup host:

- For an overview of the appliance as backup host in a virtual environment:
  See “Appliance as backup host: component overview” on page 19.

- For configuration tasks:
  See “Overview of VMware tasks” on page 24.
  See “Overview of NetBackup tasks” on page 25.
- For a list of requirements and limitations:
  See “Notes on the NetBackup appliance as a VMware backup host” on page 40.
- For log files:
  See “NetBackup logging for VMware” on page 316.

NetBackup for VMware components

Table 1-1 describes the components that NetBackup for VMware uses.

Table 1-1 Components of NetBackup for VMware

Component	Description
Backup host	NetBackup for VMware uses a special host that is called a VMware backup host (formerly called the VMware backup proxy server). The backup host is a NetBackup client that performs backups on behalf of the virtual machines. The backup host must have access to the datastores of the virtual machines. The backup host is the only host on which NetBackup client software is installed. No NetBackup client software is required on the VMware virtual machines. Note that the backup host is referred to as the recovery host when it performs a restore. The backup host can be configured in any of the following ways: ■ As a NetBackup client (Windows or Linux) with a connection to separate primary and media servers. The primary servers and media servers can be Windows, UNIX, or Linux. ■ As a NetBackup client that is installed on the media server (see Backup media server). ■ The NetBackup client and primary and media server can all reside on the same host (Windows or Linux). For a list of supported platforms for the backup host, see the NetBackup Compatibility List for all Versions.
Discovery host	Used for the automatic selection of virtual machines for backup. This host discovers virtual machines and filters them by means of the selection rules in the policy Query builder. The resulting list determines which virtual machines are backed up. The discovery host can be on any platform that NetBackup supports for primary or media servers. It can also be the same host as the backup host. You specify this host on the policy **Clients** tab: Click **Select automatically through VMware Intelligent Policy query**, then **NetBackup host to perform automatic virtual machine selection**.

Table 1-1 Components of NetBackup for VMware *(continued)*

Component	Description
Access host	A VMware Access host is another term for a NetBackup client that acts as a backup host or a recovery host.
Backup media server	A media server that can operate as the backup host. See "Media servers as backup or discovery hosts" on page 19.
NetBackup client	Installed on the backup host.
NetBackup primary server	Manages the backups of virtual machines, by means of the NetBackup client that is installed on the backup host.
NetBackup media server	Performs the backups to storage on behalf of the NetBackup client.
Virtual machine	Virtual machines provide complete guest operating systems on virtualized hardware. In a NetBackup policy, a virtual machine is configured as a NetBackup client, even though NetBackup client software is not installed on the virtual machine.
ESX server	The VMware ESX server presents a virtualized hardware environment to multiple virtual machines; each virtual machine runs an independent operating system. Users can run applications in the virtualized OS as if the OS was installed in its own physical computer.
vCenter Server	The VMware vCenter Server (or VirtualCenter server) coordinates multiple ESX servers and workloads. It can migrate virtual machines from one ESX server to another. It also provides the ability to back up the virtual machines that are turned off. The vCenter Server is optional in the NetBackup for VMware environment.

Figure 1-1 shows a NetBackup for VMware environment on a local network. The backup host accesses the VMware datastore through the ESX servers.

Figure 1-1 NetBackup for VMware: components on local network

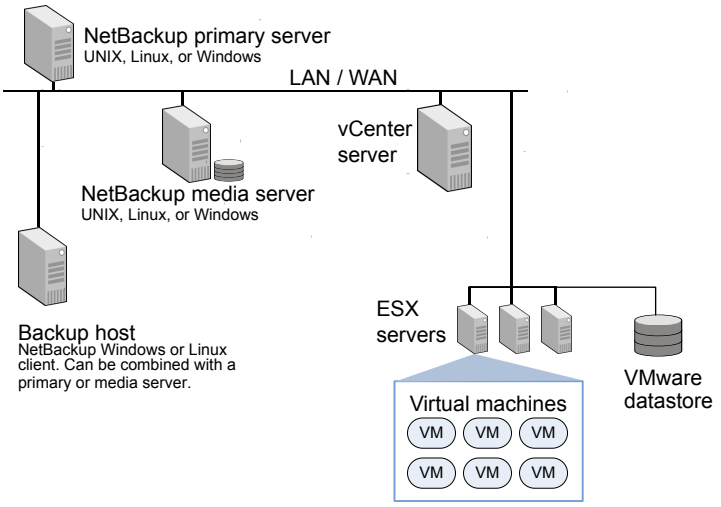

Figure 1-2 shows a NetBackup for VMware environment on a SAN. The backup host accesses the VMware datastore directly over the SAN.

Figure 1-2 NetBackup for VMware: components on SAN

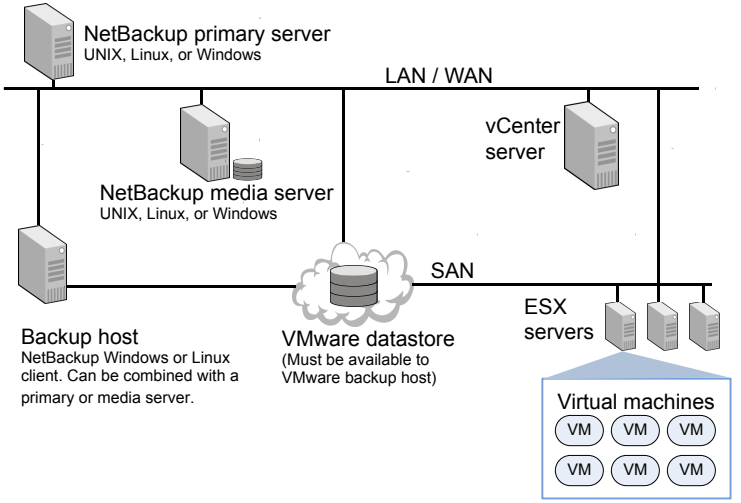

Appliance as backup host: component overview

As Figure 1-3 shows, the appliance can operate as the VMware backup host. A separate Windows backup host is not required.

The appliance as backup host can also run the NetBackup media server and primary server.

Figure 1-3 NetBackup for VMware with appliance as backup host

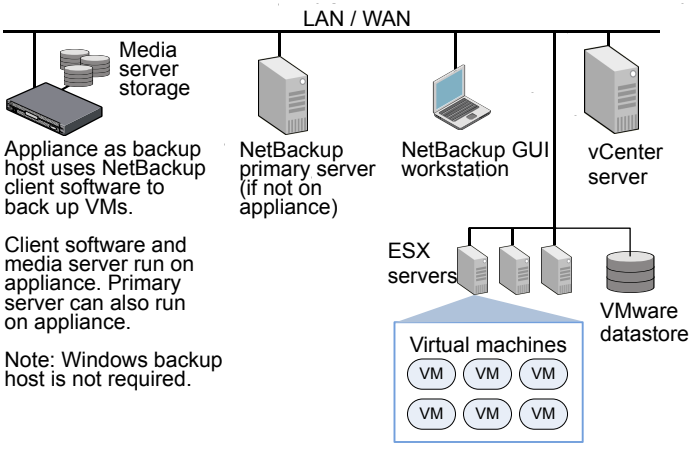

The NetBackup environment can also be on a SAN:

See Figure 1-2 on page 18.

Further information is available on the appliance as backup host:

See “Notes on the NetBackup appliance as a VMware backup host” on page 40.

Media servers as backup or discovery hosts

NetBackup for VMware uses a special host that is called a VMware backup host. The backup host is a NetBackup client that performs off-host backups of the virtual machines. The backup host must have access to the datastores of the virtual machines. The backup host reads the data from the datastore and sends it over the network to the media server. The media server backs up the data to storage.

NetBackup also uses a discovery host. For the policies that automatically select virtual machines, the discovery host filters virtual machines according to the rules in the policy Query builder. The discovery host returns a list of virtual machines to be selected for backup.

NetBackup can use media servers as backup hosts and as discovery hosts.

Media servers acting as backup or discovery hosts can provide the following advantages:

- Host redundancy: If one media server goes down, another media server takes over.
- Faster backup: The media server can read the data from the datastore and send the data straight to the storage device. Without media server access to storage devices, an ordinary backup host must send the backup data over the local network to the media server.

Figure 1-4 shows a group of media servers that can also act as backup or discovery hosts. The media servers can discover virtual machines for automatic selection, and perform off-host backups and send the backup data directly to storage.

Figure 1-4 Backup media servers

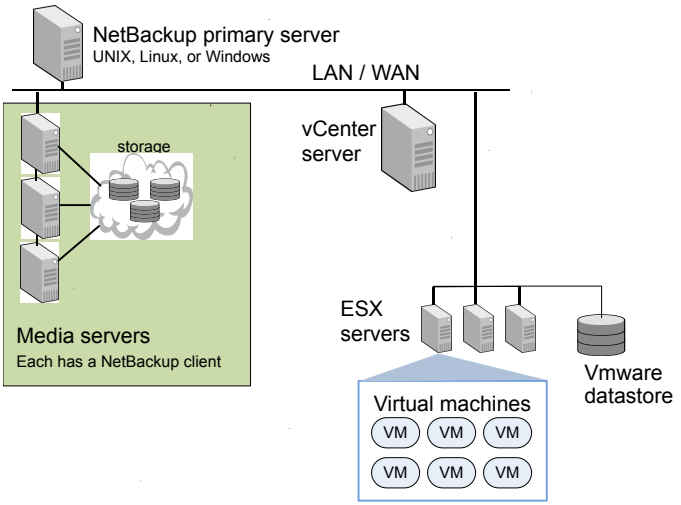

You can combine the flexibility of backup media servers with a standard feature of NetBackup: storage unit groups. Create a storage unit group that contains the storage units that your media servers can access. Any of the media servers can operate as a backup host.

Note the following requirements:

- To configure media servers as backup hosts: On the **VMware tab** of the policy set the **VMware backup host** to **Backup media server**.
  See "VMware backup host" on page 91.

- To configure media servers as discovery hosts: Set the policy **NetBackup host to perform automatic virtual machine selection** field on the **Clients** tab to **Backup media server**.
  See "Options for selecting VMware virtual machines" on page 121.

Overview of the VMware backup process

The following table describes the phases in the NetBackup backup process.

Table 1-2 NetBackup backup process

Phase	Description
Phase 1	The NetBackup primary server initiates the backup.
Phase 2	The NetBackup client on the VMware backup host initiates a VMware snapshot on the virtual machine.
Phase 3	Windows: VSS synchronizes the file system on the virtual machine. Linux: If snapshot quiesce is enabled in the Linux guest OS, the file system is synchronized on the virtual machine. (Contact your operating system vendor and VMware for additional information on how to enable snapshot quiesce.)
Phase 4	The VMware server creates a snapshot on the virtual disk datastore.
Phase 5	The NetBackup client reads the snapshot from the datastores and writes the data to the NetBackup storage unit.

NetBackup for VMware terminology

Table 1-3 lists the terminology that is used in NetBackup for VMware.

For further explanations of VMware terminology, refer to your VMware documentation.

Table 1-3 NetBackup for VMware terms

Term	Definition
backup host	The backup host is a NetBackup client that performs backups on behalf of the virtual machines. (This host was formerly known as the VMware backup proxy server.) The backup host is the only host on which NetBackup client software is installed. As an option, the backup host can also be configured as a NetBackup primary server or media server. The backup host is referred to as the recovery host when it performs a restore.
backup media server	A media server that operates as a backup host.
datastore	In NetBackup for VMware, the datastore is a disk that contains the virtual machines files.
datastore cluster	A collection of datastores that can be managed as a single unit. VMware Storage DRS manages the storage resources of the cluster.
discovery host	Discovers the virtual machines, filters them by the rules in the Query builder, and returns a list of virtual machines to be selected for backup. The discovery host is used only for automatic selection of virtual machines. Can be the same host as the VMware backup host.
query	The combination of rules in the policy's Query builder, by which NetBackup selects virtual machines for backup. A query consists of one or more rules.
Query builder	For creating filtering rules for automatic selection of virtual machines for backup. The Query builder is on the **Clients** tab of the NetBackup policy.
query rule	A single statement in a query, by which NetBackup selects virtual machines for backup. An example of a query rule is: `Displayname Contains "finance"`
recovery host	See backup host.
virtual network	A logical network that allows the exchange of data between virtual machines. A virtual network uses a virtual switch (VMware vSwitch). A virtual network can be connected to a physical network.

Table 1-3 NetBackup for VMware terms *(continued)*

Term	Definition
vmdk file	In a VMware ESX server, one or more vmdk files make up the disk image or virtual drive in a virtual machine. The `.vmdk` files contain the operating system, applications, and data in the virtual machine.
VMware Tools	Installed inside each VMware virtual machine. Enhances the virtual machine performance and adds backup-related functionality.
vmx datastore	Sometimes called the vmx directory or configuration datastore. Contains the configuration files that describe the virtual machine, such as vmx files. During a backup of a virtual machine snapshot, vmdk writes are also cached on this datastore. Note that a separate vmx datastore is not a VMware requirement.
vStorage	VMware vStorage APIs enable data protection features for more efficient use of storage capacity. NetBackup can use vStorage to back up the latest vSphere environments as well as to back up earlier VMware environments.

Chapter 2

Required tasks: overview

This chapter includes the following topics:

Overview of VMware tasks

The VMware components including ESX servers and virtual machines must be set up before you configure NetBackup.

Table 2-1 VMware tasks

Sequence	Tasks
Phase 1	Ensure that the hardware and the SAN are configured properly. The VMware datastore where the target virtual machine files exist must be accessible to the VMware backup host. A SAN connection between the backup host and the datastore is optional if you use the NBD transfer type or NBDSSL transfer type. To use the SAN transport type, set up the datastore on Fibre Channel or iSCSI. In this configuration, the VMware backup host must be able to access the datastore over the SAN. To use the hotadd transfer type for backup or restore, the VMware backup or restore host is installed in a virtual machine. See "Notes on the hotadd transport mode" on page 45.
Phase 2	Install the VMware ESX server and virtual machines.

Table 2-1 VMware tasks *(continued)*

Sequence	Tasks
Phase 3	Install VMware Tools on the virtual machines that you plan to back up. VMware requires that ESX server names resolve to an IP address. It is highly recommended that you use DNS for the naming resolution among VMware servers.
Phase 4	Optional: install a vCenter (or VirtualCenter) server.

Overview of NetBackup tasks

Table 2-2 lists the NetBackup configuration tasks for VMware. These tasks are described in other NetBackup topics and guides, as indicated.

Table 2-2 NetBackup tasks

Sequence	Tasks
Phase 1	Verify the operating system and platform compatibility. For information on supported VMware versions and on supported platforms for the backup host, see the NetBackup Software Compatibility List. The supported hardware types for the backup host are the same as for any NetBackup client. For additional support information on NetBackup for VMware, see Support for NetBackup in virtual environments.
Phase 2	Install the NetBackup primary server and media server. See the *NetBackup Installation Guide*. It is recommended that the NetBackup media server and the VMware backup host be installed on the same host.
Phase 3	Install the NetBackup Enterprise Client license on the primary server, and install NetBackup client software on the VMware backup host. NetBackup for VMware requires an Enterprise Client license for each ESX Server. To protect an application or database, note: an additional license for the appropriate NetBackup package is needed for each ESX server that hosts the application or database.
Phase 4	Install the NetBackup Java Runtime Environment (JRE) on the discovery host and the recovery host. Install the NetBackup Remote Administration Console to install NetBackup Java.

Table 2-2 NetBackup tasks *(continued)*

Sequence	Tasks
Phase 5	(Conditional) If the NetBackup primary server does not have access to the VMware server, then you need to configure a VMware backup host that has access to your NetBackup configuration. This host is also used for asset discovery. See "Add a VMware access host" on page 80. **Note:** This step is not required on the appliance: the backup host is already installed on the appliance.
Phase 6	Enter NetBackup access credentials as needed: for vCenter, for ESXi, for Restore ESXi, and for VMware Cloud Director servers. See "Add VMware servers" on page 66. **Note:** The NetBackup primary server must have network access to the VMware servers that NetBackup has credentials for.
Phase 7	Configure the NetBackup RBAC roles for VMware administrators. You must complete the discovery of VMware assets before you configure the roles. Contact your NetBackup administrator for assistance. See "RBAC roles for the VMware administrator" on page 27.
Phase 8	Create a NetBackup protection plan or policy for VMware. See "Protect VMs or intelligent VM groups" on page 173. See "Configure a VMware policy" on page 87.
Phase 9	Perform a backup. See "Manually back up virtual machines" on page 220.

Chapter 3

Configuring RBAC roles for VMware administrators

This chapter includes the following topics:

RBAC roles for the VMware administrator

NetBackup enables control over which users can access which VMware resources using Role Based Access Control (RBAC). You can grant RBAC access globally (to all VMware assets), at the individual VMware server level, or based on specific objects in the VMware object hierarchies.

The Default VMware Administrator role has access to all VMware assets (global). With this role the administrator can also manage credentials for a vCenter, ESX server, etc. (These credentials are managed on the **VMware servers** tab in **Workloads > VMware**.)

In addition, you may need other custom roles to give additional access to your VMware administrators.

- A role that gives a VMware administrator access to a guest VM credential. This way, the user can perform an agentless files and folder recovery to the guest VM without having the VM's username and password.
  See "Provide access to a credential for agentless single file recovery to a guest VM" on page 243.
- A role that is restricted to a single datacenter in a vCenter.
  See "Create a custom role for a VMware server or datacenter" on page 29.
- A role to manage an Organization VDC (OrgVDC).
  See "Create a custom role for an Organization VDC administrator" on page 30.
- A role that is restricted to an individual VM or VMs.
  See "Create a custom role to manage specific VMs" on page 31.

Note the following:

- To create an RBAC role, you must have the RBAC Administrator role or the permissions to create roles.
- To create a credential, you must have the RBAC Administrator role or a role that has permissions to create credentials. The **Default VMware Administrator** role can assign a credential to a user, but cannot create a credential in credential management.
- Contact your NetBackup administrator for assistance with creating roles and credentials.
- An RBAC role can be configured with access only on a vCloud Director server or on objects in the vCloud Director hierarchy. Users with this role are not able to see jobs in the Activity monitor.

Assigning permissions at specific VMware object levels

NetBackup supports two VMware object hierarchies, the vSphere (vCenter) object hierarchy and VMware Cloud Director (VCD) object hierarchy. Access to the applicable virtual machines is granted automatically, based on the VCD object hierarchy under which they exist.

The RBAC administrator can also assign permissions at different object levels.

You can assign permissions at the following levels:

- VMware vSphere server (vCenter)
- Datacenter in vSphere (vCenter)
- Individual VM in vCloud Director or vSphere (vCenter)

- VMware Cloud Director server
- Organization in vCloud Director
- Organization virtual datacenter (OrgVDC) in vCloud Director
- Individual vApp in vCloud Director

To manage permissions

1. On the left, go to **Workloads > VMware**. Then select the **VMware servers** tab.
2. On the left, select the object in the hierarchy.
3. On the right, locate the object and select **Actions > Manage permissions**.
4. Select the **Add** button.
5. Select the role name and the permissions that you want to assign.
6. Select the **Save** button.

Create a custom role for a VMware server or datacenter

A custom role can allow the administrator to manage a specific VMware server or datacenter. Use this role if you do not want users to have the Default VMware Administrator role. See the following topic for requirements to create an RBAC role.

See “RBAC roles for the VMware administrator” on page 27.

To create a custom role for a VMware server or datacenter

1. On the left, select **Security > RBAC** and select **Add**.
2. Select **Default VMware Administrator** and select **Next**.
3. Provide a **Role name** and a description.

   For example, include a description that the role allows users to manage the name of a specific VMware server or datacenter.
4. Under **Workloads**, select **Edit**.
5. Clear the option **Apply permissions to all existing and future VMware assets**.
6. Select **Add**.
7. Expand **VMware servers**. Then locate and select the VMware server name.

8 In the right pane, select the datacenter. Then select **Add**.

9 Select **Assign**.

10 Under **Users**, select **Edit**. Then add the users that you want to have this RBAC role.

11 Select **Assign**.

12 When you are done configuring the role, select **Add role**.

Create a custom role for an Organization VDC administrator

A custom role can allow the administrator to manage an Organization VDC. Use this role if you do not want users to have the Default VMware Administrator role. See the following topic for requirements to create an RBAC role.

See "RBAC roles for the VMware administrator" on page 27.

To create a custom role for an Organization VDC administrator

1 On the left, select **Security > RBAC** and select **Add**.

2 Select **Default VMware Administrator** and select **Next**.

3 Provide a **Role name** and a description.

For example, include a description that the role allows users to manage the name of a specific Organization VDC.

4 Under **Workloads**, select **Edit**.

5 Clear the option **Apply permissions to all existing and future VMware assets**.

6 Select **Add**.

7 On the left, expand the correct VMware Cloud Director server.

8 Locate and select the VMware Cloud Organization.

9 In the right pane, select the VMware Cloud Organization VDC. Then select the **Add** button.

10 Select **Assign**.

11 Under **Users**, select **Edit**. Then add the users that you want to have this RBAC role.

12 Select **Assign**.

13 When you are done configuring the role, select **Add role**.

Create a custom role to manage specific VMs

A custom role can allow the administrator to manage specific VMs. You can choose to manage permissions for the VMs from the VMware Cloud Director or the VMware vSphere (vCenter) object hierarchies. Use this role if you do not want users to have the Default VMware Administrator role. See the following topic for requirements to create an RBAC role.

See "RBAC roles for the VMware administrator" on page 27.

To create a custom role to manage specific VMs

1. On the left, select **Security > RBAC** and select **Add**.
2. Select **Default VMware Administrator** and select **Next**.
3. Provide a **Role name** and a description.

   For example, include a description that the role allows users to manage the names of specific VMs.
4. Under **Workloads**, select **Edit**.
5. Clear the option **Apply permissions to all existing and future VMware assets**.
6. Select **Add**.
7. On the left, expand the tree hierarchy. Then on the right pane select the VMs to which you want to grant access.
8. Select **Assign**.
9. Under **Users**, select **Edit**. Then add the users that you want to have this RBAC role.
10. Select **Assign**.
11. When you are done configuring the role, select **Add role**.

    Continue with the following steps to add additional RBAC permissions to the VMware server the RBAC role users can perform restores.
12. On the left, select **Workloads > VMware**.
13. Select the **VMware servers** tab.
14. On the right, locate the vCenter or the Datacenter where the virtual machines reside. Then select **Actions > Manage permissions**.
15. Select **Add**.
16. From the list select the new role that you created in step 1.

17 From the **Permissions** list, select the permission **View restore targets**.

18 Select **Save**.

Manage permissions for a datacenter

You can manage the RBAC permissions for a specific datacenter.

To manage permissions for a datacenter

1 On the left, select **Workloads > VMware**.

2 Select the **VMware servers** tab.

3 Select the VM and select **Manage permissions**.

4 Expand **VMware servers**. Then locate and select the VMware server name.

5 In the right pane, select the datacenter. Then select **Manage permissions**.

6 Select the role names that you want to apply to the datacenter. Or, select **Add** to add a role to the list.

To view the permissions that are applied to any role, expand the role name.

7 Select **Save**.

Manage permissions for a single VM

For a VMware VM, you can manage the RBAC permissions using vSphere object hierarchies. If a VM is managed with VMware Cloud Director (VCD), then you can also manage its permissions using the VCD object hierarchy.

To manage permissions for a VM

1 On the left, select **Workloads > VMware**.

2 Select the **Virtual machines** tab.

3 Select the VM and select **Manage permissions**.

4 Choose how you want to manage access to the VM.

- For a vSphere only VM, access is granted using the vSphere (vCenter) object hierarchy.
- For a VCD VM, you can grant access using either the VMware Cloud Director object hierarchy or the vSphere (vCenter) hierarchy. The default option is **VMware Cloud Director**.

5 Select the role names that you want to apply to the VM. Or, select **Add** to add a role to the list.

To view the permissions that are applied to any role, expand the role name.

6 Select **Save**.

Apply RBAC role permissions for a VM to other VMs

You can apply the role permissions that are applied to a VM to other VMs in the environment. Note the following requirements and limitations:

- Only the RBAC roles that are assigned at the VM asset-level can be applied to other VM assets. RBAC roles that are inherited from the VM's parent cannot be applied to other VM assets (for example, the vCenter level or the global VMware level).
- RBAC roles that are inherited from other assets cannot be applied to individual VMware assets.
- The user must have the permissions "Manage access" and "View" on a VM and a target VM to be able to view, select, and apply RBAC roles to other VMs. If a user has the **Default VMware Adminstrator** role on a VM, then they have this permission.

To apply RBAC role permissions for a VM to other VMs

1 On the left, click **Workloads > VMware**.

2 Select the VM that has the roles that you want to apply to the other VMs.

3 Select **Manage permissions**.

4 Select the role names that you want to apply to the target VMs.

To view the permissions that are applied to any role, expand the role name.

5 Click **Assign role to assets**.

6 Choose how you want to apply the selected role permissions.

- **Add to existing permissions**
  This option adds any RBAC permissions from the selected roles that do not already exist for the target VMs. The existing permissions are also retained for the target VMs.
- **Replace existing permissions**
  This option replaces the RBAC permissions for the target VMs. The existing permissions are replaced with the permissions from the selected roles.

7 Select the target VMs to which you want to apply the selected role permissions.

8 Click **Apply**.

Chapter 4

Notes and prerequisites

This chapter includes the following topics:

NetBackup for VMware: notes and restrictions

Note the following about NetBackup for VMware:

- NetBackup for VMware supports FIPS mode for SSL communication when using VDDK. This can be enabled by adding the `VDDK_FIPS_MODE = ENABLED` entry to the `bp.conf` file on VMware access hosts.
- NetBackup for VMware does not support the **Retain snapshot for Instant Recovery or SLP management** option on the policy **Attributes** tab.

  NetBackup supports these features as follows:

  - For Instant Recovery, use the `nbrestorevm` command.
    See "About Instant Recovery for VMware" on page 268.
  - For SLP management of snapshots, use Replication Director.
    For more information, see the *NetBackup Replication Director Solutions Guide*.
- NetBackup for VMware cannot back up the data on an independent disk.
  See "Troubleshooting VMware backups" on page 323.
- NetBackup for VMware does not back up standard iSCSI LUNs that are connected to the virtual machine. If the virtual machine has an iSCSI LUN, the backup succeeds but the drive that represents the LUN is not backed up.
  Note: NetBackup for VMware supports datastores over iSCSI.
- Several notes and limitations apply to Linux virtual machines.
  See "NetBackup for VMware: notes on Linux virtual machines" on page 39.
- VMware virtual machine templates are for cloning virtual machines: They cannot be turned on and used as functioning VMs. As a result, VMware has imposed the following restrictions on backup and restore of virtual machine templates:

  - A virtual machine template cannot be captured in a snapshot. NetBackup backs up the template to the designated storage unit.
  - Block level incremental backup (BLIB) cannot be used when backing up a virtual machine template. As a result of this restriction, NetBackup Accelerator cannot be used to back up VMware virtual machine templates.
  - A virtual machine template cannot be backed up over a SAN. You must configure the NetBackup policy to use a local network transfer type, such as **nbd**. Likewise, the restore of a template must be conducted over a local network. You can use the Query Builder in the NetBackup policy to create rules for automatic selection of virtual machine templates.
- NetBackup supports non-ASCII characters in virtual machine objects. Examples of objects are file and folder names, annotations, floppy image name, parallel port or serial port file name, CD-ROM ISO name, and so on.

  Support for these characters is as follows:

  - The backup host and the restore host may be the same computer.

- Windows operating systems on the NetBackup primary server, the backup host, and the restore host do not assume non-ASCII characters in VM display name if the system locales of Windows hosts are not set to UTF-8. Even if the system locale is not set to UTF-8, a backup of the VMware virtual machine whose display name contains any non-ASCII characters may work. However, a restore from the VMware virtual machine backup that has non-ASCII characters in its display name needs a Linux or UNIX restore host which uses UTF-8 character encoding.
- The UNIX and Linux operating systems on the NetBackup primary server, the backup host, and the restore host must use UTF-8 character encoding.
- The name that NetBackup uses to select a VM for backups cannot contain non-ASCII characters. The **Primary VM identifier** field in the backup policy identifies the name type that NetBackup uses to select VMs. For example, if you specify the **VM display name** as the **Primary VM identifier**, the display name of each VM that you back up cannot contain non-ASCII characters.
  See "Primary VM identifier options (VMware)" on page 93.
  VM BIOS UUID and **VM instance UUID** names never contain non-ASCII characters.
- For security purposes, VM names have some restrictions.
  See "NetBackup character restrictions for the Primary VM identifier" on page 42.

- For dual-boot virtual machines, NetBackup does not support the following policy options: **Enable file recovery from VM backup**, **Exclude deleted blocks**, **Exclude swap and paging files**, **Exclude boot disk**, **Exclude all data disks**.
- NetBackup does not support the following exclude disks options for Replication Director backups: **Exclude boot disk**, **Exclude all data disks**, **Perform custom attribute based exclusion**, **Specific disks to be excluded**.
- To back up a virtual machine while Storage vMotion migrates its files, NetBackup must conduct the backup through the vCenter server.
  See "Conflict between NetBackup and VMware Storage vMotion with vSphere 5.0 or later" on page 344.
- Several notes and limitations apply to the automatic selection of virtual machines for backup (Virtual Machine Intelligent Policy).
  See "NetBackup requirements for automatic virtual machine selection" on page 120.
- Several notes and limitations apply to the hotadd transport mode.
  See "Notes on the hotadd transport mode" on page 45.
- Note the following information about NetBackup for VMware compression and encryption:

- NetBackup's compression or encryption options
  NetBackup for VMware does not support NetBackup's compression or encryption options (in the NetBackup policy attributes).
- Granular file recovery and single file restore (SFR) on the VM
  NetBackup for VMware granular file recovery and SFR supports Windows NTFS file compression but the restored file or folder is uncompressed.
  NetBackup for VMware granular file recovery and SFR supports file-level compression (such as zip or lzh), and the file or folder is restored as the original compressed file.
  NetBackup for VMware granular file recovery and SFR does not support Windows NTFS file encryption nor any type of encryption that is set in the guest OS (such as BitLocker).
- VM volume recovery
  NetBackup for VMware VM volume recovery supports any type of compression or encryption that they are set in the guest OS (such as BitLocker).

- For limitations for the Exchange, SharePoint, and SQL Server applications, see the respective guides for those workloads.
- Several notes and limitations apply to the VMware restores.
  See "Restore notes and restrictions" on page 226.
- In vSphere, if your virtual environment has IPv6 addresses, use only the fully qualified domain names (FQDNs) that are mapped to IPv6 addresses on the DNS server.

Notes on VMware Virtual Volumes (VVols)

NetBackup supports backup and restore of the virtual machines that are configured on Virtual Volumes (VVols).

- Ensure that you have the required snapshot license from the array vendor.
- Consult the storage array documentation from the vendor for space requirements.

Configuring backup and restore of virtual machines with VVols is the same as for virtual machines without VVols, with these exceptions:

- To restore a virtual machine with the hotadd transport mode: VMware requires that the virtual machine and the restore host virtual machine reside on the same VVol datastore. Otherwise, the restore must use a different transport mode (not hotadd).
- For a restore to standard (non-VVol) datastores, the NetBackup job creates a vSphere snapshot of the virtual machine while NetBackup restores the data.

Note: For a restore to a VVol datastore, NetBackup restores the data to the virtual machine without creating a vSphere snapshot.

NetBackup IPv6 parameter required for backups in VMware IPv6 environments

For backups and restores of VMware virtual machines in an IPv6 environment, you must configure the IPv6 support on the following NetBackup hosts:

- The primary server
- The backup host

The **Both IPv4 and IPv6** option of the **Use the IP address family support** host property configures IPv6 support. This host property is located in the **Network settings** host properties.

If the NetBackup primary server and the backup host are the same host, configure the support on that host only. If the NetBackup primary server and the backup host are separate hosts, configure the support on each host.

For more information see the *NetBackup Web UI Administrator's Guide*.

NetBackup for VMware: notes on Linux virtual machines

The following notes apply to virtual machines with Linux guest operating systems:

- NetBackup cannot exclude unused or deleted blocks from the backup if the virtual machine is configured with software RAID volumes. The policy's **Exclude deleted blocks** option is not supported.
- Unmounted LVM2 volumes must start with /dev
  If the path of an unmounted LVM2 volume does not start with /dev, the backup of the virtual machine fails. Note: The path of the volume is set with the `dir` parameter on the LVM volume configuration file. An example of this configuration file is /etc/lvm/lvm.conf.
- For Linux files or directories, NetBackup for VMware has the same path name restriction as NetBackup on a Linux physical host. Files or directories with path names longer than 1023 characters cannot be individually backed up or restored. Such files can be restored when you restore the entire virtual machine from a full virtual machine backup.
- The Linux ext4 file system includes a persistent pre-allocation feature, to guarantee disk space for files without padding the allocated space with zeros.

When NetBackup restores a pre-allocated file (to any supported ext file system), the file loses its preallocation and is restored as a sparse file. The restored file is only as large as the last byte that was written to the original file. Subsequent writes to the restored file may be non-contiguous.

Note: The restored file contains all of its original data.

- The NetBackup policy's **Enable file recovery from VM backup** option is not supported for the disks inside a Linux guest OS that are configured as follows:
  - The disks are divided into logical volumes by means of the Linux Logical Volume Manager (LVM), and
  - The LVM volumes were created with thin-provisioning.
- See "Restore notes and restrictions on Linux" on page 229.

Notes on the NetBackup appliance as a VMware backup host

Note the following requirements and limitations for the appliance as the backup host:

- The appliance must be version 2.5 or later.
- You must use the VMware policy type.
- The appliance supports iSCSI connections. Refer to the *NetBackup Appliance iSCSI Guide* for more information.

NetBackup for VMware support for SAN multi-pathing

NetBackup for VMware on Windows supports dynamic multi-pathing (DMP) between the vSphere ESXi storage and the NetBackup for VMware agent (Windows backup host). Dynamic multi-pathing can provide SAN I/O high-availability and improved backup throughput.

For NetBackup for VMware on a Linux backup host, the following items describe the support for dynamic multi-pathing in a SAN environment:

NetBackup appliance

The NetBackup appliance (beginning with the 2.6.0.2 release) supports SAN dynamic multi-pathing for VMware backups. I/O is redirected through the volume manager dynamic multi-pathing node.

For more information, see the *NetBackup Appliance Administrator's Guide* for version 2.6.0.2 and later:

Non-appliance Linux host

NetBackup supports backups and restores for multi-pathing under the following conditions:

- The virtual disk SAN transport allowed list specifies the device node paths that you want to use for multipathing. The virtual disk `vixDiskLib.transport.san.allowList` API function defines the allowed list. For NetBackup purposes, Cohesity recommends that you include the DMP nodes in the allowed list. The following is an example:

  ```
  vixDiskLib.transport.san.allowList = /dev/vx/dmp/
  hitachi_usp-v0_00a0,/dev/vx/dmp/hitachi_usp-v0_00
  a0s1,/dev/vx/dmp/hitachi_usp-v0_00a0s2,/dev/vx/dm
  p/hitachi_usp-v0_00a0s3,/dev/vx/dmp/hitachi_usp-v
  0_00a0s4,/dev/vx/dmp/hitachi_usp-v0_00a0s5
  ```

- The virtual disk SAN transport blocked list specifies the device node paths you want to exclude from multipathing. The virtual disk `vixDiskLib.transport.san.denyList` API function defines the blocked list. For NetBackup purposes, Cohesity recommends that you set the blocked list to `all`. The following is an example:

  ```
  vixDiskLib.transport.san.denyList = all
  ```

Specify the allowed list and the blocked list in the following file on the backup host:

```
/usr/openv/lib/shared/vddk/lib64/vixDiskLib.ini
```

Note: For more information about the allowed list and the blocked list, see the appropriate VMware documentation. For example, for Virtual Disk Development Kit (VDDK) 6.0, the allowed list and the blocked list function descriptions are in the VMware *Virtual Disk Programming Guide*. Note that earlier versions of VDDK refer to the allowed list as the `whitelist` and the blocked list as the `blacklist`.

Note: This support or limitation does not affect NetBackup support for VMware's Native Multipathing, which is multi-pathing between the ESXi host and storage.

This support or limitation does not affect NetBackup for VMware on Windows backup hosts. The Windows operating system has integrated multi-pathing support. The following Microsoft guide contains more information:

http://microsoft.com/mpio

NetBackup for VMware support for fault tolerant VMs

NetBackup supports backing up and restoring primary fault tolerant virtual machines on vSphere 6.0 and later. (NetBackup does not backup or restore the secondary VMs.) If you choose to overwrite the VM during the restore, NetBackup deletes both the primary and the secondary VMs during the restore process. NetBackup then restores the primary VM.

Fault tolerance is not enabled on the restored VM. After the restore has completed, you can turn on fault tolerance.

For information about how to turn on fault tolerance, see the VMware documentation for your version of vSphere.

NetBackup character restrictions for the Primary VM identifier

For VMware virtual machines in a NetBackup policy, certain characters are not allowed in their names. The backup policy **Primary VM identifier** field identifies the name type that NetBackup uses to select VMs.

See "Primary VM identifier options (VMware)" on page 93.

If the name contains disallowed characters, backups or restores may fail.

The following table describes the characters and strings that NetBackup does not allow for backup or restore in the **Primary VM identifier**, except where noted.

Table 4-1 Disallowed characters and strings in the primary VM identifier

Character/string	Description	Notes
"	Quotation mark, unicode x22.	
$	Dollar sign, unicode x24.	

Table 4-1 Disallowed characters and strings in the primary VM identifier *(continued)*

Character/string	Description	Notes
'	Apostrophe, unicode x27.	
*	Asterisk, unicode x2A.	
,	Comma, unicode x2C.	
:	Colon, unicode x3A.	
;	Semi-colon, unicode x3B.	
?	Question mark, unicode x3F.	
@	At sign, unicode x40.	
|	Vertical line, unicode x7C.	
`	Grave accent, unicode x60.	
´	Acute accent, unicode xB4.	
%	Percent sign, unicode x25.	Disallowed in the virtual machine display name for application-aware VMware backups for SharePoint, allowed elsewhere. In VIP query and VM search results, NetBackup converts % in the display name of **Included** VMs to the literal string **%25**. When you specify a display name in a query, replace the % character with **%25**.
&	Ampersand sign, unicode x26.	
<	Less than sign, unicode x3C.	
>	Greater than sign, unicode x3E.	
-	Hyphen-minus, unicode x2D.	Disallowed in the first position only.
/	Solidus, unicode x2F.	Disallowed in VM display name, allowed in other objects.
\	Reverse solidus, unicode x5C.	Disallowed in VM display name, allowed in other objects.

Table 4-1 Disallowed characters and strings in the primary VM identifier *(continued)*

Character/string	Description	Notes
.	Full stop (period), unicode x2E.	Disallowed in VM display name when in the last position, allowed in other objects.
□	Space, unicode x20.	In VIP query and VM search results, NetBackup converts a space character in the display name of **Included** VMs to the literal string **%20**. When you specify a display name in a query, replace the space character with **%20**.
Unicode characters greater than x7F (non-ASCII)		Disallowed when the backup policy specifies the VM display name as the primary VM identifier.
%2f	Literal string, not a unicode character definition.	Disallowed in VM display name, allowed in other objects.
%5c	Literal string, not a unicode character definition.	Disallowed in VM display name, allowed in other objects.

Additional character restrictions for VM names can be found in the *NetBackup Cloud Administrator's Guide*, available from this location:

http://www.veritas.com/docs/000003214

In the policy Query builder, display names, resource pool names, and vApp names are case-sensitive

In VMware vSphere, virtual machine display names, resource pool names, and vApp names are case-sensitive. For example, a virtual machine with the name "vm1" is a different virtual machine from one that is named "VM1."

The VMware virtual machines with the name values that do not exactly match the query string case are not returned in the result set. Backups for those virtual machines are missed.

For example, for the following virtual machines:

```
vmware-ted
VMware-charles
VMWARE-john
vmWARE-jason
```

A query specifying `Displayname Contains "vmware-"` returns `vmware-ted` but not the other virtual machines.

Notes on the hotadd transport mode

NetBackup supports several transport modes for sending snapshot data between the VMware datastore and the VMware backup host during a backup or restore. One of those transport modes (**hotadd**) is used when the VMware backup host is installed in a virtual machine.

Note the following about the hotadd transport mode:

- The VMware backup host must be installed in a virtual machine.
- The following is a VMware requirement: The virtual machine to back up (or restore) and the virtual machine that contains the hotadd backup host must reside in the same VMware data center. The same VMware requirement applies to virtual machine restore: The virtual machine to restore and the virtual machine that contains the hotadd restore host must reside in the same VMware datacenter.
  For hotadd backup, it is recommended at least one hotadd backup host for each datacenter.
- NetBackup does not support IDE disks on the virtual machine.
- On the virtual machine to back up, no two disks should have the same name. (Identical names can occur if the disks reside on different datastores.)
- The ESX server (where the backup-host virtual machine resides) must have access to the datastore of the virtual machines that you want to back up.
- The datastore for the backup-host virtual machine must have some free space before the hotadd backup begins. Otherwise, the backup may fail.
- Locking timeouts in the VMware VDDK may cause simultaneous hotadd backups from the same VMware backup host to fail.
  See "Simultaneous hotadd backups (from the same VMware backup host) fail with status 13" on page 351.
- For a list of VMware restrictions on the hotadd transport mode, refer to VMware's documentation.

Notes and limitations for tag usage in VMware Intelligent Policy queries

- NetBackup does not support the selection of virtual machines based on the category.
- NetBackup uses tags for virtual machine selection independent of the tag's category. The VMware vSphere Web Client can create tags in two different categories with the same tag name. In the example, both virtual machines are selected if the policy is configured to include virtual machines with the “HR” tag.
  Example:
  - `Virtual_Machine_1` has a user-specified tag `HR` in the category `Production`
  - `Virtual_Machine_2` has a user-specified tag `HR` in the category `Test`
- NetBackup only recognizes the tags that are associated with virtual machines, not other vCenter objects such as Datastores.
- vCenter Server 6.0 or later is required.
- Any Windows host with only the NetBackup Client Software installed that is also defined in the VMware policy's **Client** tab as the **NetBackup host to perform automatic virtual machine selection** must have NetBackup Java installed. Install the NetBackup Remote Administration Console to install NetBackup Java.
- Be aware of a known bug in vSphere 6.0 when the system time of the discovery host and the vCenter Server are not synchronized. This issue is known to cause backups to fail with a NetBackup Status Code 4263.
  VMware knowledge base article: http://kb.vmware.com/kb/2125193
- If you use block-level incremental backups (BLIBs) with a VMware VIP policy, understand the effect of changing the VMware discovery host setting. After this policy change, subsequent incremental backups will back up the full data, not only the changed data, because the policy manager loses the backup references. If the discovery host must be changed, a full schedule backup is needed after the policy change to allow subsequent incremental backups to back up the changed data.
- In large VMware environments, consider increasing the **Maximum bearer token lifetime** from the default value. Cohesity has observed issues with discovery jobs timing out because the default value is too small. The **Maximum bearer token lifetime** is a vCenter Server setting. More information on this issue is available.
  See “Troubleshooting VMware tag usage” on page 352.

Notes and limitations for the backup and restore of VMware tag associations

- Tag associations are part of the metadata of the virtual machine. NetBackup considers virtual machine tag association protection a best effort backup. Any tag collection errors are shown in the Activity Monitor for the virtual machine snapshot jobs.

  If for any reason NetBackup is unable to back up tag associations, the job completes successfully with a NetBackup Status Code 0. Any failures to retrieve tag associations in the backup are, however, reported in the job details in the Activity Monitor.
- vCenter Server 6.0 or later is required.
- VMware tag associations are only backed up when you use VMware Intelligent Policies.
- NetBackup backs up tag associations with virtual machines. NetBackup does not back up tag associations with other vCenter Server objects, such as datastores or folders.
- Tag associations for all virtual machines are retrieved during the discovery job. The tag associations are, however, stored in the backup image for each individual virtual machine. Be aware there is a time difference between the discovery job time and the backup job. If tag associations are changed between these times, the changes are not backed up.
- If tag associations of a virtual machine were successfully backed up, then NetBackup attempts to recreate tag associations with the restored virtual machine. If the recovery host version is incompatible, NetBackup completes the recovery of the virtual machine and sets the job status to 0. The recovery job details, however, provide information on the failure to recreate the tag associations.
- If tag associations of a virtual machine were not backed up successfully, then NetBackup does not attempt to recreate tag associations with the restored Virtual Machine. NetBackup completes the recovery of the virtual machine and sets the job status to 0. The recovery job details provide information on the failure to collect the tag associations that occurred at the time of backup.
- When Replication Director for VMware is used to protect your virtual machines, tag associations are backed up only when Application Consistent protection is enabled. When Application Consistent protection is disabled, NetBackup does not protect tag associations.

Notes and limitations for the backup and restore of VMware storage policies

- NetBackup supports protecting and restoring distinct storage policies which are assigned to a VM's home directory and also to any virtual disk.
- A vCenter VM can have one or more storage policies assigned to it. NetBackup captures storage policy information during a VM backup and can apply the same or a different storage policy on a full-VM restore.
- A best effort attempt to capture a VM's storage policy information is made during the backup process.
- If a VM's storage policy information was captured during a backup and **Use a storage Policy to select datastore** option is selected during the recovery, the default storage policy information displayed is the information captured during backup.
- If a VMs storage policy information was not captured during backup, NetBackup attempts to restore the selected storage policy assigned to a VM's home directory and/or its virtual disks during recovery of the VM. If the storage policy cannot be applied, NetBackup completes the recovery of the virtual machine to the selected datastore and set the job status to 1.
- To restore a VM with storage policies that are different than what was captured during backup time, the web UI or CLI needs to be used for the restore.
- To troubleshoot storage policy-related issues during backup, refer to the `bpfis` logs. To troubleshoot storage policy-related issues during restore, refer to the `brpd` and `bpVMutil` logs.
- vMotion derived restore of an Instant Recovery (IR) or Instant Access (IA) VM, does not apply storage policies.
- To apply the storage policies information on full VM restore, the recovery host version must be 10.3 or later.
- vMotion derived restore of an Instant Recovery (or Instant Access) VM does not apply storage policies.

Support for LVM thin pool based volumes

NetBackup support the indexing of files on XFS formatted volumes and partitions for VMware Replication Director and Integrated Snapshot Manager for VMware.

Support added for LVM2 Thin Pool Volumes for VMware.

Note: Ensure to provide -V option while creating LVM thin volumes when you create thin-pool based volumes. If user fails, then the Index From Snapshot (IFS) job fails unexpectedly. For more details, refer the `lvcreate` manual page for creating LVM thin volumes.

Note: Windows backup host is not supported for the indexing of files on LVM2 thin volumes for VMware policy and VMware Replication Director and Integrated Snapshot Manager for VMware.

Chapter 5

VMware vSphere privileges

This chapter includes the following topics:

About VMware vSphere privileges

This chapter covers the privileges that NetBackup requires to work with the VMware vSphere and vCenter. These VMware vSphere privileges are granted from the VMware vSphere Web UI. They are granted to a role, and you need to use the credential of a user with that role in NetBackup when you add a VMware server. This credential is referred as “server credential” in the NetBackup documentation.

VMware vSphere privileges for virtual machine backups

This topic details the privileges that are required for virtual machine backups with various VMware vSphere transport modes. Each mode has specific privileges necessary for operation within the vSphere infrastructure. These privileges need to be set at the vCenter server level.

Transport mode: NBD

Global

- Disable methods
- Enable methods

Virtual machine

- Change configuration
  - Toggle disk change tracking
- Provisioning
  - Allow read-only disk access
  - Allow virtual machine download
- Snapshot management
  - Create snapshot
  - Remove snapshot

Transport mode: hotadd

Additional privileges are for required for the hotadd transport mode.

Datastore

- Browse datastore
- Low level file operations

Virtual machine

Change configuration

- Add existing disk
- Remove disk

Provisioning

- Allow read-only disk access

- Allow virtual machine download

Transport mode: SAN

Additional privileges are for required for the SAN transport mode.

Virtual machine

Change configuration

- Acquire disk lease

Provisioning

- Allow read-only disk access
- Allow virtual machine download

VMware vSphere privileges for a full VM restore

This section details the required privileges for different transport modes during a full VM restore in VMware vSphere. Each transport mode has the specific privileges that must be allocated to ensure a successful operation. These privileges need to be set at the vCenter server level.

Privileges for transport mode: NBD

Datastore

- Allocate space
- Browse datastore
- Low level file operations

Global

- Enable methods
- Disable methods

Network

- Assign network

Resource

- Assign virtual machine to resource pool

Virtual machine

- Change configuration
  - Acquire disk lease
  - Add existing disk

  - Add new disk
  - Add or remove device
  - Advanced configuration
  - Change settings
  - Change Swapfile placement
  - Change resource
  - Remove disk
  - Toggle disk change tracking
- Edit inventory
  - Create new
  - Remove
- Provisioning
  - Allow disk access
  - Allow read-only access
  - Allow virtual machine download
- Snapshot management
  - Create snapshot
  - Remove snapshot
  - Revert to snapshot

Privileges for transport mode: hotadd

Additional privileges are for required for the hotadd transport mode.

Datastore

- Update virtual machine files
- Update virtual machine metadata

Privileges for transport mode: SAN

The SAN transport mode requires the same privileges as NBD.

VMware vSphere privileges to create an instant access VM

This section outlines the necessary privileges create an instant access VM in VMware vSphere. Privileges are categorized based on different roles and components within the vSphere environment. These privileges need to be set at the vCenter server level.

Create an instant access VM

Datastore

- Allocate space

Global

- Enable methods
- Disable methods

Host

- Configuration
- Storage partition configuration

Network

- Assign network

Resource

- Assign virtual machine to resource pool

Virtual machine

- Change configuration
  - Add existing disk
  - Add new disk
  - Advanced configuration
  - Change Swapfile placement
  - Toggle disk change tracking
- Edit inventory
  - Create new
  - Unregister

Download files and folders

No vSphere privileges are needed for this specific operation, as all operations are performed within the NetBackup application. However, a user must be able to log into the vCenter and have the necessary setup for instant access.

VMware vSphere privileges for NetBackup plug-in operations

This section details the privileges necessary for NetBackup plug-in operations within VMware vSphere, focusing on an instant virtual machine recovery and a full VM restore with different transport modes. These privileges need to be set at the vCenter server level.

See the section called “Instant virtual machine recovery” on page 55.

See the section called “Full VM restore with transport mode: NBD” on page 57.

See the section called “Full VM restore with transport mode: hotadd” on page 58.

See the section called “Full VM restore with transport mode: SAN” on page 58.

Instant virtual machine recovery

For instant virtual machine recovery, a set of privileges are required to perform the operations that are related to the datastore, global settings, host configuration, and more.

Datastore

- Allocate space
- Browse datastore
- Low level file operations

Global

- Enable methods
- Disable methods

Host

- Configuration
  - Storage partition configuration

Network

- Assign network

Resource

- Assign virtual machine to resource pool

Sessions

- Validate session

Virtual machine

- Change configuration
  - Acquire disk lease
  - Add existing disk
  - Add new disk
  - Add or remove device
  - Advanced configuration
  - Change settings
  - Change Swapfile placement
  - Change resource
  - Remove disk
  - Rename
  - Toggle disk change tracking
- Edit inventory
  - Create new
  - Register
  - Remove
  - Unregister
- Interaction
  - Power off
  - Power on
- Provisioning
  - Allow disk access
- Snapshot management
  - Create snapshot
  - Remove snapshot

Full VM restore with transport mode: NBD

The NBD transport mode requires additional privileges for the datastore, global settings, host configuration, and more.

Datastore

- Allocate space
- Browse datastore
- Low level file operations

Global

- Enable methods
- Disable methods

Network

- Assign network

Resource

- Assign virtual machine to resource pool

Sessions

- Validate session

Virtual machine

- Change configuration
  - Acquire disk lease
  - Add existing disk
  - Add new disk
  - Add or remove device
  - Advanced configuration
  - Change settings
  - Change Swapfile placement
  - Change resource
  - Configure Raw device
  - Remove disk
  - Rename
  - Toggle disk change tracking

- Edit Inventory
  - Create new
  - Remove
- Interaction
  - Power on
- Provisioning
  - Allow disk access
- Snapshot management
  - Create snapshot
  - Remove snapshot
  - Revert to snapshot

Full VM restore with transport mode: hotadd

The hotadd transport mode requires additional privileges to update virtual machine files and allow read-only disk access and VM downloads.

Datastore

- Update virtual machine files
- Update virtual machine metadata

Virtual machine

- Provisioning
  - Allow read-only disk access
  - Allow virtual machine download

Full VM restore with transport mode: SAN

The SAN transport mode requires the same privileges as the NBD transport mode.

VMware vSphere privileges for instant rollback

This section details the privileges necessary for instant rollback operations. These privileges need to be set at the vCenter server level.

Global

- Disable methods
- Enable methods

Virtual machine

- Change configuration
  - Toggle disk change tracking
- Interaction
  - Power off
  - Power on
- Provisioning
  - Allow disk access
  - Allow read-only disk access
  - Allow virtual machine download
- Snapshot management
  - Create snapshot
  - Remove snapshot

VMware vSphere privileges for agentless SFR privileges

This section outlines the privileges that are required for agentless Single File Restore (SFR) using different transport modes in VMware vSphere: NBD, hotadd, and SAN. The privileges that are indicated are specifically for restoring to the original location of the virtual machine (VM). These privileges need to be set at the vCenter server level.

Agentless SFR transport modes: NBD, hotadd, SAN

The following privileges are necessary across all transport modes to perform agentless SFR operations within VMware vSphere.

Global

- Enable methods
- Disable methods

Resource

- Assign virtual machine to resource pool

Virtual machine

- Change configuration

  - Add existing disk
  - Add new disk
  - Add or remove device
  - Remove disk
  - Toggle disk change tracking
- Edit inventory
  - Create new
  - Remove
- Guest operations
  - Guest operation modifications
  - Guest operation program execution
  - Guest operation queries
- Provisioning
  - Allow disk access
  - Allow read-only access
  - Allow virtual machine download
- Snapshot management
  - Create snapshot
  - Remove snapshot

VMware vSphere privileges for individual vmdk restore privileges

This section details the privileges that are required for individual virtual machine disk (vmdk) restore operations within VMware vSphere. The privileges that are indicated are specifically for restoring to the original location of the virtual machine (VM). These privileges need to be set at the vCenter server level.

Global

- Enable methods
- Disable methods

Resource

- Assign virtual machine to resource pool

Virtual machine

- Change configuration
  - Add existing disk
  - Add new disk
  - Add or remove disk
  - Remove disk
  - Toggle disk change tracking
- Edit inventory
  - Create new
  - Remove
- Interaction
  - Power off
  - Power on
- Provisioning
  - Allow disk access
  - Allow read-only access
  - Allow virtual machine download
- Snapshot management
  - Create snapshot
  - Remove snapshot
  - Remove snapshot

VMware vSphere privileges for vApp restore and vApp restore to template

This section lists the required privileges for vApp restore and vApp restore to template operations in the VMware vSphere.

These privileges need to be set at the vCenter server level.

Datastore privileges

- Allocate space

- Browse datastore

Global privileges

- Enable methods
- Disable methods

Network privileges

- Assign network

Resource privileges

- Assign virtual machine to resource pool

vApp privileges

- Add virtual machine
- Assign resource pool
- Create
- Move
- Power off
- Power on

Virtual machine privileges

- Change configuration
  - Acquire disk lease
  - Add existing disk
  - Add new disk
  - Add or remove device
  - Advanced configuration
  - Change settings
  - Change Swapfile placement
  - Modify device settings
  - Remove disk
  - Toggle disk change tracking
- Edit inventory
  - Create new
  - Register

- Provisioning
  - Allow disk access
  - Allow read-only access
- Snapshot management
  - Create snapshot
  - Remove snapshot
  - Revert to snapshot

Optional permissions for better integration with VMware vSphere

These permissions allow NetBackup to send backup-related events and create and set custom attributes or annotations. The NetBackup plug-in for VMware vSphere Client can take advantage of this data but is not required to see it from the vCenter Web Client UI.

Global

- LogEvent
- SetCustomAttribute

Note: If encryption policies are used, additional privileges for cryptographic operations may be necessary.

The **Post vCenter events** option in the advanced attributes of a VMware policy lets NetBackup send backup-related events to the vCenter server.

See "Post vCenter events option (VMware advanced attributes)" on page 101.

Chapter 6

Managing VMware servers

This chapter includes the following topics:

About VMware discovery

NetBackup automatically starts the discovery of the VMware server when you add a VMware server or update credentials. The backup host information is used to validate the credentials and perform the discovery. Discovery occurs at set intervals. (The default interval is every 8 hours.) Starting with NetBackup 10.5, the version

of a backup host must be at NetBackup 10.5 or later. It is recommended that you use a backup host whose version is the same as the primary server.

After an upgrade to NetBackup 10.5 from a previous release, you must manually run discovery to refresh the list of VMware assets. This refresh lets you select specific VMware datacenters when you configure an RBAC role.

You can change the autodiscovery frequency.

See "Change the autodiscovery frequency of VMware assets" on page 65.

To discover the VMs immediately:

See "Discover VMware server assets manually" on page 65.

Change the autodiscovery frequency of VMware assets

Automatic discovery of VMware assets occurs at regular intervals. The default frequency is every 8 hours. Use this procedure to change the autodiscovery frequency.

To change the frequency of autodiscovery of VM assets

1. On the left, click **Workloads > VMware**, then click the **Virtual machines** tab.
2. On the right, select **VMware settings > Autodiscovery**.
3. Select **Frequency > Edit**.
4. Use the up or down arrows to choose how often you want NetBackup to perform autodiscovery of VMware assets. Then click **Save**.

   The range from which you may choose is 1 hour to 24 hours. To set the autodiscovery frequency in minutes or seconds or to disable autodiscovery, you must use the VMware autodiscovery API.

Discover VMware server assets manually

Use this procedure to manually discover any VMware server so that you can view and protect recently added assets.

Note: Automatic discovery of VMs and other objects in the vCenter, ESXi server, VMware Cloud Director server, or Restore ESXi server begins: when server credentials are added or updated through the web UI or an API. However, the server's VMs and other objects might not appear in the UI immediately. They appear after the discovery process for the VMware server completes. Discovery also occurs at set intervals according to the `VMWARE_AUTODISCOVERY_INTERVAL` option. (The default interval is every 8 hours.) More information about this option is available:

See "Change the autodiscovery frequency of VMware assets" on page 65.

To manually discover VMware server assets

1. On the left, click **Workloads > VMware**, then click the **VMware servers** tab.

   The tab lists: the names and types of vCenters, standalone ESXi servers, VMware Cloud Director servers, and Restore ESXi servers that you have access to. You can also review the **Discovery Status** and **Last discovery attempt** to determine when the server's VMs and other objects were last discovered.

2. Locate and select the VMware server.

3. Select **Actions > Discover**.

   The discovery operation may fail if the VMware server credentials are invalid. To validate and update the credentials:

   See "Validate and update VMware server credentials" on page 71.

For more information about the protection status of VMs and intelligent VM groups:

See "View the protection status of VMs or intelligent VM groups" on page 178.

See "Troubleshooting the status for a newly discovered VM" on page 330.

Add VMware servers

NetBackup requires logon credentials for a VMware server for either of the following reasons:

- To browse the server's virtual machines and back them up.
- To use the server as a target for restoring virtual machines.

See the section called "Notes on server names" on page 70.

To add VMware servers and their credentials

1. On the left, click **Workloads > VMware**, then click the **VMware servers** tab.

   The tab shows the servers that you can access.

2. Click **Add** to add a server.

3. Select the server type.

   See the section called “Server types and their credentials” on page 67.

4. Enter the host name.

   See the section called “Notes on server names” on page 70.

5. Add the credentials.

6. Choose a **Backup host for validation**.

   Starting with NetBackup 10.5, the version of a backup host must be at NetBackup 10.5 or later. It is recommended that you use a backup host whose version is the same as the primary server.

   See “Add a VMware access host” on page 80.

7. Indicate a **Port** number for the connection.

   If the default port number has not been changed on the VMware server, no port specification is required. If the VMware server has been configured to use a different port, specify that port number.

8. Click **Save**.

   VMs and other objects appear after the discovery process for the VMware server completes.

Server types and their credentials

Table 6-1 describes the types of VMware servers that you can add and any requirements when you add their credentials.

Table 6-1 Server types

Server type	Description
vCenter	Designates a vCenter (or VirtualCenter) server that manages ESX servers. When you create a policy to back up this server's virtual machines, NetBackup can browse this server and list its virtual machines. If the credentials provide full access privileges to the vCenter server, you can restore virtual machines to this server. Note: Do not enter logon credentials for the individual ESX servers that this vCenter server manages. NetBackup needs credentials for the vCenter only. If you enter credentials for both an ESXi server and a vCenter that manages it, problems such as the following may occur: ■ A VMware Intelligent Policy (VIP) may fail due to discovery of duplicate VMs. ■ For manual selection of VMs (not VIP): Tags and custom attributes (which require a vCenter) may cause intermittent problems if NetBackup attempts to back up the VM using the ESXi server.
ESXi	Designates a standalone ESXi server that a vCenter server does not manage. NetBackup can browse the ESXi server to present a list of its virtual machines for backup. You can also restore virtual machines to this ESXi server. To use the server as a target for restores, enter the credentials that provide full access privileges to the ESXi server.

Table 6-1 Server types *(continued)*

Server type	Description
Restore ESXi	Designates an ESXi server to which NetBackup can restore virtual machines. You must enter the credentials that provide full access privileges to the server. **Note:** NetBackup accesses this type of server for restores only, not for backups. The restore ESXi server type has the following advantages: ■ For large environments with hundreds of hosts, NetBackup may not need full access to the vCenter server. With the restore ESXi server type, you can give NetBackup full access to a single ESXi server that is dedicated to restore. ■ SAN-based restores that go directly to a restore ESXi server are faster than restores that go through the vCenter server. ■ Allows restoring to an ESXi 5.x or later server that a vCenter 5.x or later server manages. NetBackup uses vCenter to create the virtual machine. NetBackup then writes the .vmdk files directly to the ESXi server using the Restore ESXi Server credentials to that server. **Note:** VMware does not support the restore of virtual machines directly to an ESXi 5.x or later server that vCenter manages. To restore the virtual machine, select the vCenter server as the destination. As an alternative, you can set up an independent ESXi server to be used for restores. You must add NetBackup restore credentials for that ESXi server by means of the **Restore ESXi** type. For further information on the restore ESX server, refer to the following Cohesity tech note: http://www.veritas.com/docs/000007351
VMware Cloud Director	Designates a Cloud Director server. NetBackup can browse the vCloud environment on this server to present a list of its virtual machines for backup. You can also restore virtual machines to this server. Note the following: ■ The credentials must be for a system administrator account. ■ For backup and restore to Cloud Director, both vCloud and vCenter credentials are required (**vCenter**). ■ If the vCloud environment uses a load-balancer to manage multiple cells (nodes), add credentials for the load balancer, not for the cells. If Cloud Director has multiple cells but no load balancer, add credentials for only one of the cells, not for all of them. Note also: If the Domain Name System (DNS) cannot resolve the name of the load balancer or cell, do the following: Include a line in the hosts file on the VMware backup host that specifies the IP address of the load balancer or cell.

Notes on server names

Refer to the following guidelines when you enter the server name.

- Enter the server name in the same format in which it is registered in DNS and in the VMware server (whether short or fully-qualified).
  See “Using the VMware Managed Object Browser to verify the server name” on page 70.
  Fully qualified names are recommended. The entire name must be properly formed without empty or null elements. For example, a fully-qualified name must include the domain name and not end in a period (.).
- For the vCenter name, note the following:
  - The vCenter name is case-sensitive.
  - The vCenter name must match the name that is set on the vCenter for `VimApiUrl` name and `Runtime` name. For assistance in setting those names on the vCenter, and for additional vCenter naming requirements relating to the NetBackup plug-ins for vSphere: See the topic on consistent vCenter naming in the NetBackup Plug-in for VMware vSphere Client (HTML5) Guide.
- The ESX server name is case-sensitive. Enter the ESX server name exactly as it is in the VMware environment. If the case is wrong, the credential validation fails and states "...expecting <correct_name_of_server>."

Using the VMware Managed Object Browser to verify the server name

When adding NetBackup credentials for a VMware server, enter the server name exactly as configured in the VMware server (whether short or fully-qualified). If the name you enter for the credentials does not match the name as defined on the VMware server, the credential validation fails.

You can use the vSphere Managed Object Browser to verify the server's name.

To verify the server's name

1. In a web browser, open the Managed Object Browser (MOB) by entering the fully-qualified domain name of the VMware server and `/mob`.

   For example: `https://vcenter1.acmecorp.com/mob`

2. Navigate to the **ManagedObjectReference:HostSystem** and find the name value for the server.

   The object structure is site-dependent.

3. When you create NetBackup credentials for the server, enter the name value exactly as it appears in the MOB.

   See "Add VMware servers" on page 66.

   For more information on the MOB, see your vSphere documentation.

Validate and update VMware server credentials

After a VMware server is added, you can validate or update the credentials for the server.

To validate VMware credentials

1. On the left, click **Workloads > VMware**, then click the **VMware servers** tab.
2. Select one or more VMware servers, then click **Validate**.

   NetBackup verifies the current credentials for the selected VMware servers.

   If the credentials are not valid, NetBackup indicates **Invalid** under **Credentials**.

To update VMware server credentials

1. On the left, click **Workloads > VMware**, then click the **VMware servers** tab.
2. Locate the VMware server.
3. Select **Actions > Manage credentials**.
4. Update the credentials as needed.
5. Click **Save**.

Browse VMware servers

You can browse vCenter servers, standalone ESXi servers, and VMware Cloud Director servers to locate VMs and view their details. VM details include their protection plans and recovery points.

To browse VMware servers

1. On the left, click **Workloads > VMware**.
2. Click **VMware servers** to begin searching.

   The list includes: the names and types of vCenters, standalone ESXi servers, VMware Cloud Director servers, and the Restore ESXi servers that you have access to. You can also review the **Discovery Status** and **Last discovery attempt** to determine whether the server's VMs and other objects have been successfully discovered.

   To locate a server, you can enter a string in the search field.
3. Click on a server to begin drilling into it.

   You can navigate back to a higher level by clicking the up-arrow.
4. Click on a VM to view its protection status, recovery points, and restore activity.
5. Click **Add protection** to subscribe the VM to a plan.

Remove VMware servers

Use this procedure to remove VMware servers from NetBackup.

Note: If you delete a server, all virtual machines that are associated with the deleted VMware server are no longer protected. You can still recover existing backup images, but backups of VMs on this server fail.

To remove a VMware server

1. On the left, click **Workloads > VMware**, then click the **VMware servers** tab.

   The tab lists: the names and types of vCenters, standalone ESXi servers, VMware Cloud Director servers, and Restore ESXi servers that you have access to. You can also review the **Discovery Status** and **Last discovery attempt** to determine when the server's VMs and other objects were last discovered.
2. Locate the VMware server.
3. Select **Actions > Delete**.
4. If you are sure that you want to delete the VMware server, click **Delete**.

Create an intelligent VM group

You can create an intelligent VM group based on a set of filters called queries. NetBackup automatically selects virtual machines based on the queries and adds them to the group. You can then apply protection to the group. Note that an intelligent group automatically reflects changes in the VM environment and eliminates the need to manually revise the list of VMs in the group.

Note: The web UI must discover the VMs on each server before the query can select from them. If a VMware server was recently added in the web UI, its VMs may not have been discovered.

See "Change the autodiscovery frequency of VMware assets" on page 65.

To discover the VMs immediately:

See "Discover VMware server assets manually" on page 65.

Note: Intelligent VM groups are not supported for VMware Cloud Director VMs.

To create an intelligent VM group

1. On the left, click **Workloads > VMware**.
2. Click the **Intelligent VM groups** tab and then click **Add**.
3. Enter a name and description for the group.
4. Select the appropriate VMware server.
5. Perform one of the following:
   - Select **Include all VMs**.
     This option uses a default query to select all VMs that currently reside in the vCenter or ESXi for backup when the protection plan runs.
   - To select only the VMs that meet specific conditions, create your own query: Click **Add condition**.

6 To add a condition, use the drop-downs to select a keyword and operator and then enter a value.

The options are described after this procedure: Query options for creating intelligent VM groups.

Examples are also available: Example queries

To change the effect of the query, click **Condition** and click **AND** or **OR**, then select the keyword, operator, and value for the condition. For example:

You can also add sub-queries to a condition, if necessary. Click **Sub-query** and click **AND** or **OR**, then select the keyword, operator, and value for the sub-query condition. For example:

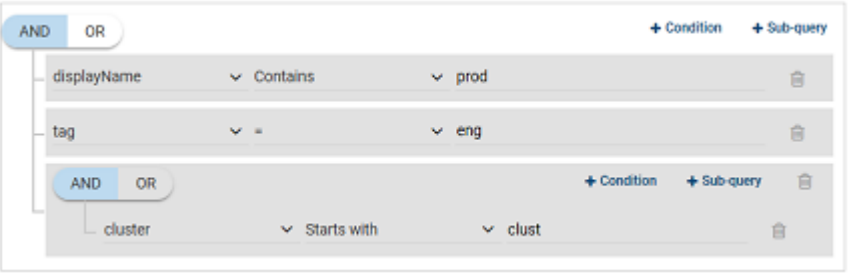

7 To test the query, click **Preview**.

The query-based selection process is dynamic. Changes in the virtual environment can affect which VMs the query selects when the protection plan runs. As a result, the VMs that the query selects later when the protection plan runs may not be identical to those currently listed in the preview.

8 To save the group without adding it to a protection plan, click **Add**.

To save and add it to a protection plan, click **Add and protect**, select the plan, and click **Protect**.

Note: When you click **Preview** or you save the group, the query options are treated as case-sensitive when the VMs are selected for the group. Under **Virtual machines**, if you click on a VM that was not selected for the group, the **Member of virtual machine groups** field reads `none`.

However, when you add the group to a protection plan, some of the query options are treated as case-insensitive when the protection plan's backup runs. As a result, the same VM may now be included in the group and is backed up.

For the case behavior of each option, see Query options for creating intelligent VM groups.

Query options for creating intelligent VM groups

Note the following for intelligent VM groups

- When using queries in **Intelligent VM groups**, the NetBackup web UI might not display an accurate list of VMs that match the query if the query condition has non-English characters. However, during the backup, the correct VMs are selected even though the VM attributes are non-English.
- Using the `not equals` filter condition on any attribute returns assets including those that have no value (null) present for the attribute. For multi-value attributes such as `tag`, the assets that do not match at least one of the values of the attribute are not returned
- When the server of an Intelligent VM group is updated, all existing access definitions configured for that Intelligent group are removed because the intelligent group is now registered with the new server namespace. You need to add new access definitions for the updated Intelligent group.

Table 6-2 Query keywords

Keyword	Description	Case-sensitive when protection plan runs
`annotation`	The text that is added to VM annotations in a vSphere client.	Yes
`connectionState`	The status of the VM connection to the ESX server. For example, if a virtual machine's ESX server is down, that virtual machine is not connected.	No
`cluster`	The name of the cluster (group of ESXi servers) where the VMs reside.	No
`datacenter`	The name of the datacenter.	No
`datacenterPath`	The folder structure that defines the path to a datacenter. Use this option if the datacenter name that you want to filter on is not unique in your environment.	Yes
`datastore`	The name of the datastore.	Yes
`displayName`	The VM's display name.	Yes
`host`	The name of the ESXi server. The ESXi host name must match the name as defined in the vCenter server.	No
`dnsName`	The VM's DNS name in vSphere Client.	No
`guestOS`	The VM guest OS type that is recorded in the vSphere client.	Yes
`hostName`	The VM name that is derived from a reverse lookup of its IP address.	No
`instanceUuid`	The VM's instance UUID. For example: `501b13c3-52de-9a06-cd9a-ecb23aa975d1`	No
`networkName`	The name of the network switch (on an ESX server) or distributed switch.	No
`powerState`	The power state of the VM.	No
`tag`	The name of the VM's tag.	Yes
`template`	Indicates if the VM is a virtual machine template.	No
`version`	The VMware version of the virtual machine. For example, vmx-04, vmx-07, vmx-08.	Yes

Table 6-2 Query keywords *(continued)*

Keyword	Description	Case-sensitive when protection plan runs
`vmFolder`	The name of the VM folder (within a datacenter), which includes the path to the folder that contains the VMs. See the section called "VMFolder examples" on page 79.	No
`vmxDatastore`	The name of the VMX datastore (sometimes called the vmx directory or configuration datastore).	Yes
`vmxDatastoreType`	The type of the VMX datastore. Values are NFS or VMFS.	No

Query operators

Table 6-3 Query operators

Operator	Description
`Starts with`	Matches the value when it occurs at the start of a string. For example: If the value you enter is `box`, this option matches the string `box_car` but not `flatbox`.
`Ends with`	Matches the value when it occurs at the end of a string. For example: If the value you enter is `dev`, this option matches the string `01dev` but not `01dev99` or `devOP`.
`Contains`	Matches the value you enter wherever that value occurs in the string. For example: If the value you enter is `dev`, this option matches strings such as `01dev`, `01dev99`, `devOP`, and `development_machine`.
=	Matches only the value that you enter. For example: If the value you enter is `VMtest27`, this option matches `VMTest27` (same case), but not `vmtest27`, `vmTEST27`, or `VMtest28`.
!=	Matches any value that is not equal to the value that you enter.

Example queries

In this example, the query adds to the group any VM that has `prod` in its display name.

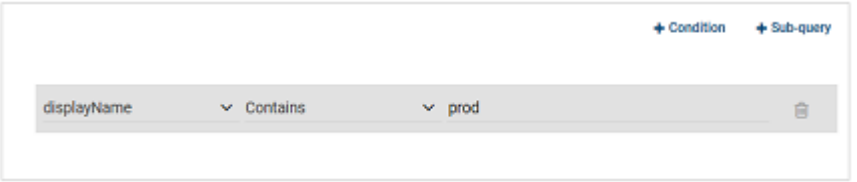

To change the effect of the query, click **Condition** and click **AND** or **OR**, then select the keyword, operator, and value for the condition. For example:

This example uses **AND** to narrow the scope of the query: it selects only the VMs that have `prod` in their display name and that also have a tag named `eng`. If a VM does not have `prod` in its display name as well as a tag named `eng`, that VM is not added to the group.

To broaden the scope of the query, use **OR**:

In this example, **OR** causes the query to add the following to the group:

- The VMs that have `prod` in their display name (regardless of any tags).
- The VMs that have a tag named `eng` (regardless of the display name).

You can also add sub-queries to a condition, if necessary. Click **Sub-query** and click **AND** or **OR**, then select the keyword, operator, and value for the sub-query condition. For example:

In this example, the sub-query causes the query to narrow the scope further. From the VMs that have both `prod` in their display name and a tag named `eng`, only the VMs in clusters that start with `clust` are selected.

VMFolder examples

For example, assume the following VM folders containing a total of 65 VMs:

`vm\VM_backup_prod1` (contains 5 VMs)

`vm\VM_backup_prod1\cluster1`(contains 10 VMs)

`vm\VM_backup_prod2` (contains 50 VMs)

To include the VMs in `vm\VM_backup_prod1` but not the VMs in `cluster1` or in any other folder:

```
VMFolder Equal "vm\VM_backup_prod1"
```

To include the VMs in `vm\VM_backup_prod1` and in its subfolder `cluster1`:

```
VMFolder Equal "vm\VM_backup_prod1"
```

OR

```
VMFolder StartsWith "vm\VM_backup_prod1"
```

Note: The first backslash is an escape character that causes the following backslash to be interpreted as a literal character.

To include all 65 VMs: `VMFolder StartsWith "vm\VM_backup_prod"`

Note: Any VM that is in a path that begins with `vm\VM_backup_prod` is included.

Remove an intelligent VM group

Use the following procedure to remove an intelligent VM group.

To delete an intelligent VM group

1. On the left, click **Workloads > VMware**.
2. Locate the group under the **Intelligent VM groups** tab.
3. If the group is not protected, select it and then click **Delete**.
4. If the group is protected, click on the group, scroll down and click the lock symbol, and click **Unsubscribe**.
5. Click **Remove**.

Add a VMware access host

NetBackup uses a special host that is called a VMware access host. This host is a NetBackup client that performs backups on behalf of the virtual machines. The access host is the only host on which NetBackup media server or client software is installed. No NetBackup client software is required on the virtual machines. However, the access host must have access to the datastores of the virtual machines. The access host reads the data from the datastore and sends it over the network to the media server.

The VMware access host was formerly called the VMware backup host or the VMware backup proxy server. The access host is referred to as the recovery host when it performs a restore.

Note the following:

- Starting with NetBackup 10.5, the version of an access host must be at NetBackup 10.5 or later. It is recommended that you use an access host whose version is the same as the primary server.
- You do not need to follow this procedure for the media servers that operate as backup hosts with the policy option **Backup media server**. With that option, NetBackup automatically enables media servers as backup hosts.
- Make sure that NetBackup media server software or client software is installed on any access host that you add.

To add a VMware access host

1. On the left, click **Workloads > VMware**, then click the **Virtual machines** tab.
2. On the right, select **VMware settings > Access hosts**.

   NetBackup lists any access hosts that were previously added.

3 Click **Add**.

4 Enter the name of the access host and then click **Add**.

5 If the NetBackup primary server is clustered in a failover environment: repeat this procedure to add the backup host to each primary server node in the cluster.

Remove a VMware access host

To remove a VMware access host

1 On the left, click **Workloads > VMware**, then click the **Virtual machines** tab.

2 On the right, select **VMware settings > Access hosts**.

NetBackup lists any access hosts that were previously added.

3 Locate the VMware access host and then click the delete icon.

4 To confirm the deletion, click **Delete**.

Change resource limits for VMware resource types

VMware resource limits control the number of backups that can be performed simultaneously on a VMware resource type. The settings apply to all NetBackup policies for the primary server that you selected.

For example, to avoid overloading the ESX server, you can place a limit on the number of concurrent backup jobs per ESX server. To control I/O overhead on the datastore's array, you can limit the number of concurrent backups per datastore.

See the section called “Limitations on global limits for VMware resources” on page 83.

To change the resource limits for VMware resource types

1 Review the limitations for resource limits.

See the section called “Limitations on global limits for VMware resources” on page 83.

2 On the left, click **Workloads > VMware**.

3 On the top right, select **VMware settings > Resource limits**.

For each resource, the default value is **0** (No limit).

Limits indicates the number of simultaneous backups that can be performed for the resource type. This value is the global limit. The **Override** value indicates how many resources have any limits that are different from the global limit.

4 Select the VMware resource type you want to change and then **Edit**.

See "VMware resource types and limits" on page 83.

Note: The **Snapshot** resource limit is different from the other resource types. It sets a limit for the number of simultaneous snapshot-only operations within a vCenter domain, such as create snapshot and delete snapshot. This limit applies only during the snapshot creation and snapshot deletion phases of a backup. It does not control the number of simultaneous backup jobs. This **Snapshot** limit can be useful for controlling the effect that multiple snapshot operations have on the vCenter server. Add a specific vCenter to override the global snapshot setting for that vCenter.

5 Choose from the following options.

Set a global limit for a VMware resource type.	Locate the **Global** setting and select the **Limits** value that you want to apply. This value limits the number of simultaneous backups that are performed for the resource type.
Set a limit for a specific VMware resource.	Click **Add**. From the list, select the resource. Select the **Limits** value that you want to apply. This value limits the number of simultaneous backups that are performed for the selected resource.

At any point you can Click **Reset default values** to remove all the overrides and set all global VMware resource limits to their default values.

6 Click **Save**.

7 To enable resource limits for restores, select the **Apply limits to restore jobs** check box.

Note: Restore resource limits are cumulative. The limit represents the number of combined backup and restores that can run simultaneously against a resource type. Restore resource limits only apply to vCenter, ESX server, and Datastore resource types.

Restores to ESXi clusters, datastore clusters, and storage policies do not count against resource limit settings. As VMware determines the placement of the VM once it is created.

Limitations on global limits for VMware resources

The following limitations apply to setting global limits on the use of VMware resources:

- New and changed resource limits may not take effect immediately. It can take a couple of jobs before the resource limit updates are in effect.
- The resource limits settings apply only to policies that use automatic selection of virtual machines (Query Builder). If virtual machines are selected manually on the Browse for Virtual Machines screen, the resource limit settings have no effect.
- To limit the number of simultaneous jobs per policy, use the **Limit jobs per policy** setting on the policy **Attributes** tab. The effect of this option depends on how the policy selects virtual machines.

VMware resource types and limits

Table 6-4 Resource types and limits

Resource type	Resource limit
vCenter	The maximum number of simultaneous backups per vCenter server.
snapshot	The maximum number of simultaneous snapshot operations (create or delete) per vCenter.
Cluster	The maximum number of simultaneous backups per VMware cluster.
ESXserver	The maximum number of simultaneous backups per ESX server.
VMXDatastore	Controls the maximum number of simultaneous backups per Datastore. The Datastore is defined as the location of the VMX file that is associated with each VM. This resource type is useful for VMs that have vmdk files distributed across multiple Datastores. This setting is enforced globally within a NetBackup domain. See "NetBackup for VMware terminology" on page 21.
Datastore	The maximum number of simultaneous backups per datastore.
DatastoreFolder	The maximum number of simultaneous backups per datastore folder.

Table 6-4 Resource types and limits *(continued)*

Resource type	Resource limit
DatastoreType	The maximum number of simultaneous backups per datastore type.
VMXDatastoreNFSHost	Controls the maximum number of simultaneous backups per Datastore at the NFS host level where the Datastore type is NFS. The **VMXDatastoreNFSHost** is an NFS server that sources one or more NFS Datastores. For this resource type the Datastore is defined as the location of the VMX file that is associated with each VM. This resource type is useful for VMs that have any vmdk files that are distributed across multiple Datastores. This setting is enforced globally within a NetBackup domain.
DatastoreNFSHost	The maximum number of simultaneous backups per NFS host of the datastore.
DatastoreCluster	The maximum number of simultaneous backups per datastore cluster.

Setting privileges for posting events to vCenter

With the **Post vCenter events** option, NetBackup can send backup related events to the vCenter server, to view in vSphere Web Client. The NetBackup plug-in for vSphere Client (HTML5) is not required.

Note the following requirements:

- You must enter the credentials that give NetBackup access to the vCenter server.
  See "Add VMware servers" on page 66.
- Make sure that the **Post vCenter events** option is enabled in the policy.
  See "VMware - Advanced attributes" on page 96.
- You must set the correct vCenter role privileges. Use the vSphere Web Client to make sure that the following **Global** privileges are set in vCenter: **Manage custom attributes** and **Set custom attribute**.
  For assistance setting privileges, refer to the appropriate VMware vSphere Documentation Center.

Authentication token for the NetBackup vSphere plug-ins

With the vSphere Client (HTML5) plug-in, the VMware administrator can use the vSphere interface to recover virtual machines. To allow the plug-in to communicate with the NetBackup primary server, you must provide an authentication token to the VMware administrator.

To create an authentication token see the NetBackup Plug-in for VMware vSphere Client (HTML5) Guide.

Validating VMware virtualization server certificates in NetBackup

NetBackup can validate VMware virtualization server certificates using their root or intermediate certificate authority (CA) certificates.

For more information on external CA support in NetBackup, refer to the NetBackup Security and Encryption Guide.

The following procedure is applicable for the NetBackup primary server and all VMware access hosts.

To configure secure communication between VMware virtualization server and VMware access host

1. Configure an external CA trust store on the VMware access host.
2. Add CA certificates of the required VMware servers (vCenter, ESX, or ESXi server) in the trust store on the access host.

   For the Windows certificate store, you need to add the CA certificate to the Windows Trusted Root Certification Authorities.

   Use the following command:

   ```
   certutil.exe -addstore -f "Root" certificate filename
   ```
3. Use the `nbsetconfig` command to configure the following NetBackup configuration options on the access host. See the *NetBackup Administrator's Guide, Volume I* for details on these options.

ECA_TRUST_STORE_PATH	Specifies the file path to the certificate bundle file that contains all trusted root CA certificates. This option is specific to file-based certificates. You should not configure this option if the Windows certificate store is used. If you have already configured this external CA option, append the VMware CA certificates to the existing external certificate trust store. If you have not configured the option, add all the required virtualization server CA certificates to the trust store and set the option.
ECA_CRL_PATH	Specifies the path to the directory where the certificate revocation lists (CRL) of the external CA are located. If the configuration option is already configured, append the virtualization server CRLs to the CRL cache. If the option is not configured, add all the required CRLs to the CRL cache and then set the option.
VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED	Lets you enable the validation of a virtualization server's certificate.
VIRTUALIZATION_CRL_CHECK	Lets you validate the revocation status of the virtualization server certificate against the CRLs. By default, the option is disabled.
VIRTUALIZATION_HOSTS_CONNECT_TIMEOUT	Lets you specify the duration (in seconds) after which the connection between NetBackup and vCloud Director server ends.
VMWARE_TLS_MINIMUM_V1_2	Lets you specify the Transport Layer Security (TLS) version to be used for communication between NetBackup and VMware servers.

Chapter 7

Configuring backup policies for VMware

This chapter includes the following topics:

Configure a VMware policy

You can create a NetBackup policy (full or incremental) to back up the virtual machine.

Note: To configure a policy for Replication Director, see the *NetBackup Replication Director Solutions Guide*.

To configure a policy to back up the virtual machine

1. On the left, select **Protection > Policies**.
2. Enter a name for the policy.
3. For **Policy type**, select **VMware**.

   VMware backup options are available on the **VMware** tab (described later in this procedure).
4. Select a policy storage unit or storage unit group.
5. In most cases, you can leave the **Disable client-side deduplication** option at the default (cleared).
6. To enable the Accelerator, select **Use Accelerator**.

   This action selects the **Enable block-level incremental backup** option on the **VMware** tab.

   See "About the NetBackup Accelerator for virtual machines" on page 160.
7. To define a schedule, click the **Schedules** tab and click **Add**.

   For assistance with the **Accelerator forced rescan** option:

   See "Accelerator forced rescan for virtual machines (schedule attribute)" on page 164.

   On the **Attributes** tab, you can select **Full backup**, **Differential incremental backup**, or **Cumulative incremental backup**.

   Note that incremental backups require one of the following selections on the **VMware** tab of the policy:

   - **Enable file recovery from VM backup**
   - **Enable block-level incremental backup**
     This option requires an ESX server 4.0 and a virtual machine at vmx-07 or later.
     More information is available for incremental backups.
     See "About incremental backups of virtual machines" on page 112.

8 Click the **Clients** tab to select the virtual machines to back up.

Note: If you change the **Virtual machines for backup** option from manual selection to intelligent policy (or vice versa), the next backup of the VM is a regular full backup. This action occurs even if a backup already exists for that VM. For a policy that uses **Enable block-level incremental backup** (BLIB) or BLIB plus Accelerator, the backup processing is not limited to changed blocks only.

The options for selecting virtual machines are as follows:

Select manually, and click **Add**.

You can type the host name in the **Enter the VM hostname** field, or select **Browse virtual machines**.

See "Browse for VMware virtual machines" on page 107.

Select automatically through VMware intelligent policy query

With this option, NetBackup can automatically select virtual machines for backup based on the filtering criteria that you enter. The following topics explain how to specify the criteria:

See "About automatic virtual machine selection for NetBackup for VMware" on page 115.

See "Configure automatic virtual machine selection" on page 125.

Enable VMware Cloud Director integration

Enables backup of the virtual machines that reside in a vCloud environment. With this option, the policy selects for backup only the virtual machines that vCloud manages: it skips the virtual machines that are not in vCloud.

See "About NetBackup for vCloud Director" on page 254.

NetBackup host to perform automatic virtual machine selection

This option displays when you click **Select automatically through VMware intelligent policy query**. This host discovers virtual machines and automatically selects them for backup based on the query rules.

See "About automatic virtual machine selection for NetBackup for VMware" on page 115.

9 Select the **VMware** tab to set VMware-related options.

See "Backup options on the VMware tab" on page 90.

10 To exclude disks from the backups, select the **Exclude disks** tab.

See "Exclude disks tab" on page 101.

11 Click **Create**.

Limit jobs per policy on the Attributes tab (for VMware)

The **Limit jobs per policy** option operates as follows, depending on how the policy selects virtual machines.

For the policies that select virtual machines automatically (Query builder)

The **Limit jobs per policy** option controls the number of parent (discovery) jobs that run simultaneously for the policy. This option does not limit the number of snapshot jobs and backup (bpbkar) jobs that the parent job launches. For example, if this option is set to 1 and you begin a backup of a policy that discovers 100 virtual machines: all the snapshot jobs and backup jobs for each of the 100 virtual machines are allowed to run simultaneously. Only the initial discovery job counts against **Limit jobs per policy**. If you begin a second backup of the policy, its discovery job cannot start until all the child jobs from the first backup are complete.

For the policies that use manual selection of virtual machines

Limit jobs per policy controls the number of virtual machines that the policy can back up simultaneously. Because no discovery job is needed, each virtual machine backup begins with a snapshot job. Each snapshot counts against the **Limit jobs per policy** setting. If this option is set to 1: the backup of the next virtual machine that is specified in the policy cannot begin until the first snapshot job and its backup are complete.

See “Change resource limits for VMware resource types” on page 81.

Backup options on the VMware tab

The **VMware** tab displays when you select VMware as the policy type.

The following topics describe the VMware backup options.

See “VMware backup host” on page 91.

See “Optimizations options (VMware)” on page 91.

See “Primary VM identifier options (VMware)” on page 93.

See “Existing snapshot handling options (VMware)” on page 94.

See “Transport modes options (VMware)” on page 95.

See “Application protection options (VMware)” on page 96.

See “VMware - Advanced attributes” on page 96.

VMware backup host

The VMware backup host is a NetBackup client that performs backups on behalf of the virtual machines.

Table 7-1 VMware backup host selection

Option	Description
Backup media server	This option allows a media server that is selected in the policy to operate as the backup host. (The storage unit determines the selection of the media server.) To operate as the backup host, the media server must contain NetBackup client software. **Note:** The storage unit that is specified in the policy must be unique to the media servers that NetBackup supports as VMware backup hosts. If the storage unit is available on a media server that is not a supported VMware backup host, the snapshot may not succeed (status 20). For a list of supported platforms for the VMware backup host, see the NetBackup Software Compatibility List (SCL) **Note:** When the **Backup media server** option is selected, NetBackup cannot determine a host to perform policy validation. To validate the policy, temporarily select one of the possible media servers as the backup host (do not select **Backup media server**). When the policy validates successfully, reset the backup host to **Backup media server**. See "Media servers as backup or discovery hosts" on page 19.
backup_host_name	Select a backup host to perform the backup. The list contains any media servers that are supported as backup host. It also contains any NetBackup clients that were added to the **VMware access hosts** list. For a list of supported platforms for the backup host, see the *NetBackup Software Compatibility List*.

Optimizations options (VMware)

The following options set the type and scope of the VMware virtual machine backup.

Table 7-2 Optimizations

Option	Description
Enable file recovery from VM backup	This option allows restore of individual files from the backup. With or without this option, you can restore the entire virtual machine. You can also use this option for incremental backups: in the policy schedule, select differential incremental backup or cumulative incremental backup. To perform a VMware backup to a deduplication storage unit, select this option. This option provides the best deduplication rates. To back up a virtual machine that contains Cohesity Storage Foundation Volume Manager volumes, disable this option. Also make sure that the **Exclude deleted blocks** option is disabled. **Note:** For a Linux virtual machine, the name of an LVM volume can include any of the following special characters: . (period), _ (underscore), - (hyphen). No other special characters are supported. If other special characters are in the volume name, the **Enable file recovery from VM backup** option does not work. As a result, you cannot restore individual files from that volume. **Note:** During an incremental backup, any files that had been moved or renamed are not backed up. Those files are not available when you browse to restore individual files from the incremental backup. However, when you restore the entire VM from a block-level incremental backup, note: the file metadata is updated and the moved or renamed files in the restored VM reflect the updated metadata.
Enable block-level incremental backup	For block-level backups of the virtual machine. This option reduces the size of the backup image. For the Exchange, SQL, and SharePoint Agents, this option is selected and grayed out if you enable **Use Accelerator**. On the **Attributes** tab, **Perform block level incremental backups** is automatically selected and grayed out.
Exclude deleted blocks	Reduces the size of the backup image by excluding any unused or deleted blocks within the file system on the virtual machine. This option supports the following file systems: Windows NTFS, and Linux ext2, ext3, ext4, and XFS. This option uses proprietary mapping technology to identify vacant sectors (allocated but empty) within the file system. To back up a virtual machine that contains Cohesity Storage Foundation Volume Manager volumes, disable this option. Also make sure that the **Enable file recovery from VM backup** option is disabled.

Table 7-2 Optimizations *(continued)*

Option	Description
Exclude swap and paging files	Reduces the size of the backup image by excluding the data in the guest OS system paging file (Windows) or the swap file (Linux). **Note:** This option does not exclude the swapping and paging files from the backup: it only excludes the data in those files. If the files are restored, they are restored as empty files. **Note:** For a Linux virtual machine, this option disables the swap file when you restore the virtual machine. You must reconfigure the swap file after the virtual machine is restored. To allow the virtual machine to be restored with its swap file enabled, do not select **Exclude swap and paging files**.

Primary VM identifier options (VMware)

This setting specifies the type of name by which NetBackup recognizes virtual machines when it selects them for backup.

The name you use for the **Primary VM identifier** may have restrictions.

See "NetBackup character restrictions for the Primary VM identifier" on page 42.

Table 7-3 Primary VM identifier

Option	Description
VM hostname	The network host name for the virtual machine. (This option is the default.) NetBackup obtains the host name by means of a reverse lookup on the virtual machine's IP address. **Note:** For NetBackup to look up the IP address, the virtual machine must already be turned on. If no host name can be found, the IPv4 address is used as the host name. In case of an IPv6 address, the field remains blank. NetBackup cannot select a VMware virtual machine for backup if it cannot obtain an IP address for the virtual machine. See "Preventing browsing delays caused by DNS problems" on page 335.

Table 7-3 Primary VM identifier *(continued)*

Option	Description
VM display name	The name of the virtual machine as displayed in the VMware interface. A display name is assigned to the virtual machine when the virtual machine is created. When virtual machines are included in a NetBackup policy, restrictions apply to the characters that are allowed in the virtual machine display name. See "NetBackup character restrictions for the Primary VM identifier" on page 42. **Note:** The restrictions also apply to other vSphere objects, such as floppy image name, parallel port or serial port file name, and CD-ROM ISO name. Each display name must be unique in your VMware environment. See "NetBackup for VMware: notes and restrictions" on page 35.
VM BIOS UUID	The ID assigned to the virtual machine when the virtual machine is created. This ID may or may not be unique, depending on whether the virtual machine has been duplicated. This option is included for compatibility with the policies that use the older VM UUID identifier.
VM DNS Name	The VMware DNS Name of the virtual machine. In vSphere Client, this name appears on the virtual machine's Summary tab. **Note:** This name may or may not be associated with the virtual machine's IP address. VMware Tools obtains this name from the host name that is configured in the virtual machine. For further information on this name, refer to the documentation for the guest operating system.
VM instance UUID	The globally unique ID assigned to the virtual machine when the virtual machine is created. This ID uniquely identifies the virtual machine within a vCenter server. Even if the virtual machine has been duplicated (such as within a vCloud), only the original virtual machine retains this instance ID. (The virtual machine duplicates are assigned different instance UUIDs.) This option is recommended instead of the **VM BIOS UUID** option.

Existing snapshot handling options (VMware)

This option specifies the action that NetBackup takes when a snapshot is discovered before NetBackup creates a new snapshot for the virtual machine backup. After it creates a snapshot, NetBackup usually deletes the snapshot when the backup completes. If snapshots are not automatically deleted (whether created by NetBackup or not), the performance of the virtual machine may eventually decline.

Undeleted snapshots can cause restore failures due to lack of disk space. If the virtual machine was configured on multiple datastores and a leftover snapshot

existed on the virtual machine when it was backed up, note: NetBackup tries to restore all .vmdk files to the snapshot datastore. As a result, the datastore may not have enough space for the .vmdk files, and the restore fails. (For a successful restore, you can restore the virtual machine to an alternate location. Select a datastore for the .vmdk files.)

Table 7-4 Existing snapshot handling: Options

Option	Description
Remove NetBackup snapshots and continue the backup	If a virtual machine snapshot exists that a NetBackup backup previously created: NetBackup removes the old snapshot, creates an updated snapshot, and proceeds with the virtual machine backup.
Continue the backup	NetBackup ignores any existing virtual machine snapshots (including snapshots previously created by NetBackup) and proceeds with snapshot creation and the backup.
Stop the backup if any snapshots exist	If any snapshot exists on the virtual machine, NetBackup stops the job for that virtual machine only.
Stop the backup if NetBackup snapshots exist	If a virtual machine snapshot exists that a NetBackup backup previously created, NetBackup stops the job for that virtual machine only.

Transport modes options (VMware)

The transport modes determine how the snapshot data travels from the VMware datastore to the VMware backup host. The appropriate mode depends in part on the type of network that connects the VMware datastore to the VMware backup host.

By default, all modes are selected. NetBackup tries each transport mode in order, from top to bottom. It uses the first mode that succeeds for all disks in the virtual machine.

Table 7-5 Transport Modes

Mode	Description
SAN	For unencrypted transfer over Fibre Channel (SAN) or iSCSI. **Note:** This mode is not supported for the virtual machines that use VMware Virtual Volumes (VVols).

Table 7-5 Transport Modes *(continued)*

Mode	Description
HotAdd	Lets you run the VMware backup host in a virtual machine. **Note:** For the virtual machines that use VVols, the virtual machine and the backup host (hotadd) virtual machine must reside on the same VVol datastore. For instructions on this transport mode and on installing the backup host in a VMware virtual machine, refer to your VMware documentation.
NBD	For unencrypted transfer over a local network that uses the Network Block Device (NBD) driver protocol. This mode of transfer is usually slower than Fibre Channel.
NBDSSL	For encrypted transfer (SSL) over a local network that uses the Network Block Device (NBD) driver protocol. This mode of transfer is usually slower than Fibre Channel.
Actions > Move up **Actions > Move down**	Use these options to change the order in which NetBackup tries each selected mode. For example: assume that all four transport modes are selected, and the order is SAN, HotAdd, NBD, and NBDSSL. If one of the virtual disks cannot be accessed using SAN, the SAN transport mode is not used for any of the virtual machine's disks. NetBackup then tries to use the HotAdd mode for all the disks. NetBackup continues to try each mode until it finds one that succeeds for all the disks.

Application protection options (VMware)

To enable file-level recovery of database data that resides in the virtual machine, select from the following options. These options apply to full backups of the virtual machine; they do not apply to incremental backups.

To configure VMware backups of database data, refer to the appropriate NetBackup database agent guide.

VMware - Advanced attributes

The following additional parameters are available for VMware backups. In most situations, the best settings are the defaults.

Table 7-6 VMware advanced attributes

Configuration parameter	Description
Virtual machine quiesce	This option is enabled by default. In the great majority of cases, you should accept the default. I/O on the virtual machine is quiesced before NetBackup creates the snapshot. Without quiescing file activity, data consistency in the snapshot cannot be guaranteed. If not consistent, the backed-up data may be of little or no value. If this option is disabled, the snapshot is created without quiescing I/O on the virtual machine. In this case, you must perform your own analysis for data consistency in the backup data. Caution: Cohesity does not recommend that you disable quiesce. In most cases, this option should be enabled. **Note:** To use this option, VMware Tools must be installed on the virtual machine. **Note:** To use this option with Linux virtual machines, snapshot quiesce must be enabled in the Linux guest OS.
Ignore Instant Recovery VMs	If this option is enabled (the default): NetBackup skips any virtual machine that was restored with Instant Recovery for VMware if the virtual machine is running from a NetBackup NFS datastore. When the virtual machine data files have been migrated to the production datastore, the virtual machine can be backed up. NetBackup identifies Instant Recovery virtual machines according to the following criteria: ■ The virtual machine has a snapshot that is named `NBU_IR_SNAPSHOT`, ■ And the virtual machine is running from a datastore and the name of the datastore begins with `NBU_IR_`. If the virtual machine meets all these criteria, it is not backed up if this option is enabled. If this option is disabled: NetBackup backs up the virtual machine even if it is running from the NetBackup NFS datastore.

Table 7-6 VMware advanced attributes *(continued)*

Configuration parameter	Description
Treat Tags as unset if unable to evaluate	Tags were introduced with VMware vCenter Version 5.1. The APIs to interface with the tagging service were not released until VMware vCenter Version 6.0. NetBackup for VMware supports the tags that are assigned to virtual machine objects starting with VMware vCenter Version 6.0. If you have a mixed vCenter environment, such as 5.1, 5.5, and 6.0: you can use this configuration parameter to modify how NetBackup treats tags for vCenter Version 5.1/5.5. This configuration parameter also applies to all versions of ESXi hosts whose credentials were added to the **Virtual Machine Servers** list. This option is disabled by default. If you use the tag field in the VMware Intelligent Policy query and your policy searches for virtual machines across a mixed vCenter environment, note: NetBackup reports the virtual machines that are discovered from vCenter Server 5.1/5.5 and ESXi hosts as failed if it needs to evaluate the tag portion of the query to make an include or exclude decision. When you enable this option NetBackup treats tags as **unset**. If you use the tag field in the VMware Intelligent Policy query and your policy searches for virtual machines across a mixed vCenter environment, note: NetBackup evaluates the tag part of the query as if no tags were set on the virtual machines that are discovered from vCenter Server 5.1, 5.5, and ESXi.
Ignore diskless VMs	If this option is enabled: NetBackup does not back up a replicated (passive) VM in a vCenter Site Recovery Manager (SRM) environment if that VM has no vmdk files. NetBackup skips that VM and backs up the corresponding active VM, which has vmdk files. Note that virtual machines without vmdk files can occur in a vCenter SRM environment. If a replicated virtual machine has never been active, it is in passive mode and may have no vmdk files. **Note:** If this option is enabled and NetBackup does not have access to the vCenter where the active virtual machine runs: the policies in the Query Builder run without error; no attempt is made to back up the virtual machine. For the policies that use manual selection of virtual machines, backups fail with status 156, because the virtual machine cannot be located. If this option is disabled: NetBackup attempts to back up a virtual machine regardless of whether it has vmdk files. If the virtual machine has no vmdk files, the backup fails with status 156.
Multiple organizations per policy	This option is disabled by default. If it is enabled, the query rules can select virtual machines from different vCloud Director organizations and back them up to the same storage unit. If you do not want backups of virtual machines from different organizations to be stored on the same drive, leave this option disabled.

Table 7-6 VMware advanced attributes *(continued)*

Configuration parameter	Description
Continue VIP discovery if one vSphere login fails	**Note:** This option applies to VMware Intelligent Policies (VIP) only. When this option is set to Yes: For a VIP policy's discovery job, NetBackup ignores a failed logon to a vCenter and attempts to log on and discover VMs on other vCenters. On any vCenter that NetBackup can log on to, the VMs that match the VIP policy's query are backed up. When this option is set to No (the default): If the attempt to log on to a vCenter fails, the discovery job fails and no VMs are backed up for any vCenters. More information is available on the types of NetBackup jobs for VMware: See "Using the Activity monitor to monitor virtual machine backups" on page 222.
Post vCenter events	Enables NetBackup to send backup related events to the vCenter server. The events appear in vSphere Client under **Home > Inventory > Hosts and Clusters, Tasks & Events** tab. See "Viewing NetBackup activity in vSphere Client (HTML5)" on page 223. The following options are available: ■ **All events**: NetBackup posts an event to the vCenter server on each backup success or failure. This setting is the default. ■ **No events**: Disables the **Post vCenter events** option. NetBackup does not post any events to the vCenter server. ■ **Error events**: NetBackup posts an event to the vCenter server only for backup failures. Further information on **Post vCenter events** is available: See "Post vCenter events option (VMware advanced attributes)" on page 101.
VMware server list	Specifies a colon-delimited list of virtual machine servers that NetBackup communicates with for this policy. In large virtual environments, you can use this list to improve backup performance: NetBackup communicates only with the servers that are in this list. For example, exclude from the list any vCenter or vCloud servers that do not contain virtual machines to be backed up by this policy. Also, if a duplicate of the virtual machine exists on a different server: the duplicate is not backed up if it resides on a server that is not included in this list. **Important:** IPv6 addresses are not supported in the **VMware server list** field. Use fully qualified domain names or host names. **Note:** Each host name must match exactly the name as configured in the NetBackup credentials. **Note:** Separate the names with a colon (:) not a comma (,). This option does not affect either of the following: the list of all possible values in the Query Builder for automatic selection of virtual machines, or browsing of virtual machines for manual selection. If the list is blank (the default), NetBackup communicates with any servers in the virtual environment.

Table 7-6 VMware advanced attributes *(continued)*

Configuration parameter	Description
VMDK compression	Used to specify the vmdk compression method. When this option is set to `none` (default) compression is not used during the backup job. NetBackup uses the preferred vmdk compression method when it opens VMDKs. Backup jobs automatically set the compression method to `none` if NetBackup is unable to read the vmdk with the preferred method.
Snapshot parameters	**Snapshot retry count** Sets the number of times the snapshot is retried. The default is 10. The range is 0 to 100. This option and the snapshot time-out and snapshot creation interval provide flexibility in the creation of snapshots. For most environments, the default values are usually best. In special circumstances, it may be helpful to adjust these settings. Example considerations are the size of the virtual machine and the processing load on the VMware server. **Snapshot timeout (minutes)** Sets a time-out period (in minutes) for completion of the snapshot. The default is 0, which means no time-out. If snapshots do not complete, set this option to a specific period to force a time-out. Consider using the snapshot creation interval to retry the snapshot at a later time. **Snapshot creation interval (seconds)** Determines the wait time (in seconds) before the snapshot is retried. The default is 10 seconds. The range is 0 to 3600. **Perform snapshot without quiescing if quiesced snapshots fail** This option is disabled by default. If it is enabled and a quiesced snapshot cannot be created, the snapshot is created without quiescing I/O on the virtual machine. The resulting snapshot is referred to as crash consistent. In this case, you must perform your own analysis for data consistency in the backed-up data. The associated snapshot job completes with a status of `0` (Success). You can configure this job to return status `1` (Partial success). Add the policy name to the `VM_SNAPSHOT_QUIESCE_STATUS` configuration setting in `bp.conf` on the media server and the backup host machines. See the NetBackup Administrator's Guide, Volume I for details. **Warning:** In most cases, Cohesity does not recommend enabling this option. NetBackup cannot guarantee that all required data has been flushed to disk when the snapshot occurs. The data that is captured in the snapshot may be incomplete. If this option is disabled, the backup fails if a quiesced snapshot cannot be created.

Post vCenter events option (VMware advanced attributes)

The **Post vCenter events** option is available on the **VMware** tab of a NetBackup policy in the **VMware advanced attributes** section.

Post vCenter events enables NetBackup to send backup related events to the vCenter server.

Note the following:

- To post events to vCenter, NetBackup must perform the backup through a vCenter server. If NetBackup accesses the ESX server directly, the backup information cannot be displayed in vSphere Client.
- You must set the required permissions in vCenter:
  See "Setting privileges for posting events to vCenter" on page 84.
- If a vSphere administrator created an attribute named NB_LAST_BACKUP of type Global, NetBackup cannot post backup events to that attribute. You must remove the NB_LAST_BACKUP attribute from vSphere. Make sure that **Post vCenter events** is set to All Events or Error Events. At the next backup, NetBackup creates a NB_LAST_BACKUP attribute of type Virtual Machine and posts events to that attribute.

Post vCenter events also records the date and time of the last successful backup of the virtual machine:

- The date and time appear in vSphere Client on the **Summary** tab as a custom attribute under **Annotations**. The attribute is labeled NB_LAST_BACKUP.
- The date and time appear in vSphere Web Client on the **Virtual Machines** display.

The events can also be viewed with the NetBackup vSphere Client (HTML5) plug-in:

See "Viewing NetBackup activity in vSphere Client (HTML5)" on page 223.

For instructions on installing and using the NetBackup plug-in, see the NetBackup Plug-in for VMware vSphere Client (HTML5) Guide.

More information is available about how to allow NetBackup to send backup-related events and create and set custom attributes or annotations.

See "Optional permissions for better integration with VMware vSphere" on page 63.

Exclude disks tab

The **Exclude disks** tab displays for policies of the **VMware** policy type. These options determine the kind of disks on the virtual machine that are excluded from the backup. These options can reduce the size of the backup, but should be used

with care. These options are intended only for the virtual machines that have multiple virtual disks.

The following options appear on the **Exclude disks** tab.

Table 7-7 Options on the **Exclude disks** tab of the policy

Option	Description
No disks excluded	Backs up all virtual disks that are configured for the virtual machine.
Exclude boot disk	The virtual machine's boot disk (for example the C drive) is not included in the backup. Any other disks (such as D) are backed up. Consider this option if you have another means of recreating the boot disk, such as a virtual machine template for boot drives. See "About the exclude disk options for virtual disk selection" on page 104. **Note:** A virtual machine that is restored from this backup cannot start. Data files are available in the restored data disks.
Exclude all data disks	The virtual machine's data disks (for example the D drive) are not included in the backup for this policy. Only the boot disk is backed up. Consider this option only if you have a different policy that backs up the data disks. See "About the exclude disk options for virtual disk selection" on page 104. **Note:** When the virtual machine is restored from the backup, the virtual machine data for the data disk may be missing or incomplete.

Table 7-7 Options on the **Exclude disks** tab of the policy *(continued)*

Option	Description
Perform custom attribute based exclusion	Exclude disks by a VMware Custom Attribute that is applied to a virtual machine. The VMware Custom Attribute identifies the disks that you want to exclude from backups. If you select this option, also enter the name of the Custom Attribute. NetBackup then excludes the disks that are defined in that attribute. The attribute must have comma-separated values of device controllers for the disks to be excluded. For example: `scsi0-0,ide0-0,sata0-0,nvme0-0` The default value is `NB_DISK_EXCLUDE_LIST`. You can use this value as the custom attribute or choose your own value. **Note:** Custom Attribute based disk exclusion requires that you enter in NetBackup the credentials for the vCenter server or servers that host the VMs. ESXi server credentials are not sufficient. See "Add VMware servers" on page 66. Your VMware administrator must use a VMware interface to apply the custom attribute to the disks that you want to exclude from the virtual server. The **Virtual Disk Exclusion** wizard of the NetBackup plug-ins for vSphere provides a method to add a Custom Attribute to a virtual machine or virtual machines. For more information, see the NetBackup Plug-in for VMware vSphere Web Client Guide or the NetBackup Plug-in for VMware vSphere Client (HTML5) Guide.
Specific disks to be excluded	Exclude a specific disk by selecting the disk controller type and device numbers that represent the virtual device node of the disk. Then click **Add**. NetBackup adds the controller ID to the list of nodes to exclude. Repeat for each disk that you want to exclude. To delete a disk from the list of disks to exclude, locate the disk controller type and device numbers, and then click **Delete**.

Note: NetBackup does not support the exclude disks options for Replication Director backups.

See "Exclude disks from backups: an example to avoid" on page 106.

See "Restoring data from the backups that excluded the boot disk or data disks" on page 106.

About the exclude disk options for virtual disk selection

The backup policy **Exclude disks** tab has options to exclude virtual disks from a backup. By default, no disks are excluded. You should use this setting in most cases.

If you want to exclude disks from a backup, the other options are **Exclude boot disk** and **Exclude all data disks**. These options are intended for the virtual machines that have multiple virtual disks. You should use these options with care.

To exclude a boot disk or data disk, note the following requirements:

- The virtual machine must have more than one disk.
- NetBackup must be able to identify the boot disk.
- The boot disk must not be part of a managed volume (Windows LDM or Linux LVM). The boot disk must be fully contained on a single disk.

The boot disk must include the following:

- The boot partition.
- The system directory (Windows system directory or Linux boot directory).

Important! The exclude disk options are meant only for the following cases:

- **Exclude boot disk**: Consider this option if you have another means of recreating the boot disk, such as a virtual machine template for boot drives. If **Exclude boot disk** is enabled, the policy does not back up the boot disk.

  Note: When the virtual machine is restored from the backup, the virtual machine data for the boot disk may be missing or incomplete.

  Note the following about **Exclude boot disk**:

  - If the virtual machine has a boot disk but has no other disks, the boot disk is backed up. It is not excluded.
  - If the virtual machine's boot disk is an independent disk, but the virtual machine has no other disks, the boot drive is backed up. The restored boot drive however contains no data, because NetBackup cannot back up the data in an independent disk.
  - If the virtual machine has a boot drive and an independent drive, the boot drive is not backed up. Only the independent drive is included in the backup. Since NetBackup cannot back up the data in an independent disk, the restored independent disk contains no data.
    Refer to the explanation on independent disks in the following topic:
    See “NetBackup for VMware terminology” on page 21.

- Adding a virtual disk and changing this option before the next backup can have unexpected results.
  See "Exclude disks from backups: an example to avoid" on page 106.

- **Exclude all data disks**: Consider this option if you have a different policy or other backup program that backs up the data disks. If **Exclude all data disks** is enabled in a policy, that policy does not back up the data disks.

  Note the following about excluding data disks:

  - If the virtual machine has only one disk (such as C:), that drive is backed up. It is not excluded.
  - If the virtual machine's boot disk is an independent disk, and the virtual machine has a separate data disk, the boot disk is backed up. The restored boot disk however contains no data, because NetBackup cannot back up the data in an independent disk.

  Note: When the virtual machine is restored from the backup, the virtual machine data for the data disk may be missing or incomplete.

- **Perform custom attribute based exclusion**: If this option is enabled, NetBackup excludes the disks that have a custom attribute from the backup. The default value for this attribute is `NB_DISK_EXCLUDE_DISK`. You can use this default attribute or change the attribute name on the **Exclude disks** tab of the backup policy. Note that this option gives the VMware administrator control over which disks are excluded.

  Note: When a virtual machine is restored from the backup, the virtual machine data for the excluded disk may be missing or incomplete.

  - The attribute on the virtual machine must be populated with comma-separated values of controller IDs for the disks to be excluded.
  - If the custom attribute is not populated or does not exist on the virtual machine, none of the disks (except independent disks) are excluded.
  - If you remove disks from the custom attribute value between the differential backups, only those files that changed since the last backup are available to restore individually. You can restore the entire virtual disk or the VM, in which case all files are restored including those you cannot restore individually. After the next full backup, you can restore any of the files individually.

- If you add disks to the custom attribute value between the differential backups, those disks are excluded from the next backup.

- **Specific disks to be excluded**: If this option is enabled, NetBackup excludes the disks that you specify. Note that this option gives the NetBackup administrator control over which disks are excluded from backups.

  Note: When a virtual machine is restored from the backup, the virtual machine data for the excluded disk may be missing or incomplete.

  - On the **Exclude disks** tab of the backup policy, you must select the option **Specific disks to be excluded** and add each control ID.
  - If the disks do not exist on the specified controller and device IDs, none of the disks (except independent disks) are excluded.
  - If you remove controllers from the exclusion list between the differential backups, only those files that changed since the last backup are available to restore. All files are available to restore after the next full backup.
  - If you add controllers to the exclusion list between the differential backups, their disks are excluded from the next backup.

Caution: The exclude disk options can have unintended consequences if these rules are not followed.

Exclude disks from backups: an example to avoid

You should use the options on the **Exclude disks** tab of the policy with care. For example, if you add a disk to the virtual machine and change the settings that exclude disks, note: The next backup may not capture the virtual machine in the state that you intended. You should back up the entire virtual machine (that is, do not exclude any disks) before you exclude a disk from future backups.

Restoring data from the backups that excluded the boot disk or data disks

If the policy's **Excludes disks** option excluded the boot disk or data disks, you can restore the backed-up data as follows:

- If **Enable file recovery from VM backup** was enabled on the backup policy: You can restore individual files from those portions of the virtual machine that the Virtual disk selection option did not exclude.
  See “Restore individual files and folders” on page 250.

- If the **Excludes disks** option was set to **Exclude boot disk**: You can restore the virtual machine and move the restored data disks to another virtual machine.

Browse for VMware virtual machines

When you configure a NetBackup policy, you can enter the virtual machine's host name manually or browse for and select it from a list.

As an alternative, NetBackup can automatically select virtual machines based on a range of criteria.

See "About automatic virtual machine selection for NetBackup for VMware" on page 115.

The following options are available:

- **Enter the VM hostname**
  The format of the name depends on your system. It may be the fully qualified name or another name, depending on your network configuration and how the name is defined in the guest OS. If NetBackup cannot find the name you enter, policy validation fails.

  Note: The type of name to enter depends on the **Primary VM identifier** setting on the **VMware** tab of the policy.

  To enter a name, make sure that **Browse virtual machines** is not selected.

- **Browse virtual machines**
  The virtual machine names that are listed may be derived from a cache file. Use of the cache file is faster than rediscovering the virtual machines on the network if your site has a large number of virtual machines.
  If NetBackup cannot obtain the IP address of the virtual machine, the host name and IP address are displayed as NONE.
  Note that virtual machine host names or display names must be unique within a primary server's policies.

Last update

This column shows the date and time of the most recent cache file that contains the names of virtual machines.

For NetBackup to access the virtual machines, note the following:

- The NetBackup primary server must have credentials for the VMware vCenter or ESX servers.
  See "Add VMware servers" on page 66.

- DNS problems may prevent or slow down discovery of the virtual machines.
- To limit the search to particular vCenter or ESX servers, you can create a `BACKUP` registry entry as an exclude list. Excluding unneeded servers can dramatically speed up the search for virtual machines.
- The browsing time out value must not be set too low.
  See "Changing the browsing timeout for virtual machine discovery" on page 337.

Limiting the VMware servers that NetBackup searches when browsing for virtual machines

As part of creating a NetBackup policy, you must specify which virtual machines to back up. One approach is to let NetBackup search the network and list all available virtual machines. However, if your VMware environment contains many VMware servers and virtual machines, it may take too long to search and list all of them. For example, consider an environment with ten vCenter servers. To back up the virtual machines on one of the ten vCenter servers, browsing virtual machines on all ten servers is unnecessary.

To speed up browsing, you can exclude particular VMware servers from the search. Then NetBackup queries only the VMware servers that are not named in the exclude list for the backup host.

Use one of the following procedures, depending on the platform of the backup host (Windows or Linux). As an alternative, NetBackup can automatically select virtual machines based on a range of criteria.

See "About automatic virtual machine selection for NetBackup for VMware" on page 115.

Limit the VMware servers that NetBackup discovers (Windows)

Use this procedure to limit the VMware servers that NetBackup discovers for each Windows backup host.

To limit the VMware servers that NetBackup discovers

1. On the Windows desktop of the backup host, click **Start > Run** and enter `regedit`.
2. Make a backup of the current registry (**File > Export**).
3. Go to **HKEY_LOCAL_MACHINE > SOFTWARE > Veritas > NetBackup > CurrentVersion > Config** and create a key called `BACKUP`.

4 Right-click in the right pane and click **New > String Value**. Enter `excludeVMservers` as the name.

5 Right-click the `excludeVMservers` name and click **Modify**.

6 In the **Edit String** dialog, enter a comma-delimited list of the VMware servers that are NOT to be queried when NetBackup browses the network. Do not enter spaces. You can enter vCenter servers and individual ESX servers.

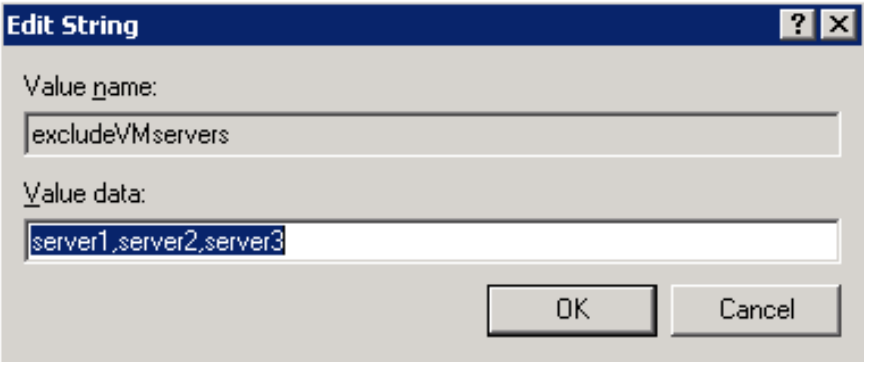

Note: The exclude list is used on the next backup. If any bpfis processes are running, the exclude list has no effect on them.

The exclude list applies only to this backup host. The servers are not queried when you manually add virtual machines to the **Clients** tab.

Limit the VMware servers that NetBackup discovers (Linux)

Use this procedure to limit the VMware servers that NetBackup discovers for each Linux backup host.

To limit the VMware servers that NetBackup discovers

1. On the Linux backup host, create (or open) the following file:

   `/usr/openv/netbackup/virtualization.conf`

2. Add the following to the file:

   ```
   [BACKUP]
   "excludeVMservers"="server1,server2,server3"
   ```

   Where *server1,server2,server3* is a comma-delimited list of the VMware servers that are NOT to be queried when NetBackup browses the network. Do not enter spaces. You can enter vCenter servers and individual ESX servers.

   Note: If the file already contains a `[BACKUP]` line, do not add another `[BACKUP]` line. Any other lines that already exist under `[BACKUP]` should remain as they are.

3. Save the file.

   Note: The exclude list is used on the next backup. If any bpfis processes are running, the exclude list has no effect on them.

   The exclude list applies only to this backup host. The servers are not queried when you manually add virtual machines to the **Clients** tab.

Virtual machine host names and display names should be unique if VMs are selected manually in the policy

Certain VMware environments do not require unique names for virtual machines. For instance, virtual machines within a vCenter server can have the same host or display names as virtual machines in another vCenter server. The same is true of datacenters, which are logical groupings of virtual resources within a vCenter server. Virtual machine host names or display names must be unique within a datacenter. They do not need to be unique between two datacenters on the same vCenter. A virtual machine named VM1 can exist in datacenter A. Another virtual machine (also named VM1) can exist in datacenter B, on the same vCenter server.

Identically named virtual machines can present a problem for any policies that are configured as follows:

- The primary server's policies use the **Select manually** option on the **Clients** tab to select the VMs for backup.
- The **Primary VM identifier** option on the **VMware** tab identifies VMs by their host names or display names.

These policies may back up a different but identically named VM, instead of the VM that you selected. In that case, the VM that you selected is not backed up. For these policies to work, the virtual machines' display names or host names must be unique.

Consider the following options:

- For manual policies that identify VMs by display name or host name, change the VM names so that each VM has a unique host name or display name.
- As an alternative, configure the policies' **Primary VM identifier** option to identify the VMs by their UUIDs instead of by host name or display name. Use the type of UUID that is appropriate for your virtual machine environment.
  See "Primary VM identifier options (VMware)" on page 93.
- Instead of policies with manual-selection, use VMware Intelligent policies to select the VMs through a query. Even if the **Primary VM identifier** option is set to host name or display name, NetBackup identifies each VM by its UUID.

Primary VM identifier option and manual selection of virtual machines

This topic describes the issues you may encounter with the **Primary VM identifier** option when you manually select of virtual machines for a policy. If the policy selects virtual machines automatically, refer to the following topic:

See "Effect of Primary VM identifier parameter on Selection column in Test Query results" on page 154.

When creating virtual machines, use the same name for both the host name and display name. If the **Primary VM identifier** is changed, the existing entries on the **Clients** tab still work. Otherwise, a change to the policy's **Primary VM identifier** value can affect backups. If you change this option, you may have to delete the virtual machine selections on the **Clients** tab and re-enter them. Then NetBackup may no longer be able to identify the virtual machines to back up.

For example, the host names in the **Clients** tab cannot be used and the virtual machines are not backed up in the following case:

- If you change the **Primary VM identifier** from **VM hostname** to **VM display name**, and
- The display names of the virtual machines are different from the host names.

In this case, delete the host name entries on the **Clients** tab and browse the network to select the virtual machines by display name.

See "Browse for VMware virtual machines" on page 107.

About incremental backups of virtual machines

NetBackup enables full virtual machine and file-level incrementals in the same backup (the **Enable file recovery from VM backup** option).

Better support for incremental backup is available when you use BLIB (**Enable block-level incremental backup**). BLIB requires ESX 4.x and virtual machines at vmx-07 or later.

Note the following:

- Individual file recovery is supported from full backups and from incremental backups, as long as the **Enable file recovery from VM backup** policy option is enabled.
- Make sure that the virtual machines to back up are time synchronized with the backup host. Otherwise, some changed data may not be included in the backup, depending on the clock differential between the backup host and the virtual machine.
- For incremental backups, you do not have to configure the client on the VMware backup host for timestamps. The VMware policies automatically default to the use of timestamps.

Configuring incremental backups

Use the following procedure for virtual machine backup.

To configure incremental backup of a virtual machine

1. In the NetBackup policy **Attributes** tab, select the **VMware** policy type.
2. On the **VMware** tab, select the VMware backup host.
3. (Optional) For BLIB, select **Enable block-level incremental backup**.
4. Select **Enable file recovery from VM backup**.

   More information is available on the VMware options.

   See “Backup options on the VMware tab” on page 90.
5. On the **Schedules** tab, select **Differential incremental backup** or **Cumulative incremental backup**.
6. Fill in the **Clients** tab.

See “About incremental backups of virtual machines” on page 112.

Storage Foundation Volume Manager volumes in the virtual machine

To back up a virtual machine that contains Cohesity Storage Foundation Volume Manager volumes, make sure the following options on the policy's VMware tab are disabled:

Enable file recovery from VM backup

Exclude deleted blocks

Note: Restore of selected files from a backup of the full virtual machine is not supported if the virtual machine contains Storage Foundation Volume Manager volumes.

See “NetBackup for VMware: notes and restrictions” on page 35.

Chapter 8

Configuring a VMware Intelligent Policy

This chapter includes the following topics:

About automatic virtual machine selection for NetBackup for VMware

Instead of manually selecting the virtual machines for backup, you can configure NetBackup to automatically select virtual machines based on a range of criteria. You specify the criteria (rules) in the **Query Builder** on the NetBackup policy **Clients** tab. NetBackup creates a list of the virtual machines that currently meet the rules and adds those virtual machines to the backup.

This feature is called the VMware Intelligent Policy.

Automatic selection of virtual machines has the following advantages:

- Simplifies the policy configuration for sites with large virtual environments. You do not need to manually select virtual machines from a long list of hosts: NetBackup selects all the virtual machines that meet the selection rules in the policy's Query Builder.
- Allows the backup list to stay up-to-date with changes in the virtual environment. Eliminates the need to revise the backup list whenever a virtual machine is added or removed.
- Virtual machine selection takes place dynamically at the time of the backup.

Examples of automatic virtual machine selection are the following:

Table 8-1 Examples for automatic virtual machine selection

Example	Description
Add new virtual machines	At the next backup, the policy can automatically discover the virtual machines that have recently been added to the environment. If the virtual machines match the query rules that you configure in the policy, they are automatically backed up.
Limit the backup list to the virtual machines that are currently turned on	If some of your virtual machines are occasionally turned off, NetBackup can be configured to automatically exclude those from the backup list. Among the virtual machines it discovers, NetBackup backs up only the virtual machines that are turned on.
Back up virtual machines based on physical boundaries	Examples of physical boundaries are vCenter servers, ESX servers, datastores, and clusters. For example, a query rule can select all the virtual machines in a particular ESX server, so the policy backs up only those virtual machines.
Back up virtual machines based on logical boundaries	Examples of logical boundaries are folders, vApps, templates, and resource pools. For example, a query rule can select all the virtual machines in a particular folder, so the policy backs up only those virtual machines.
Back up virtual machines based on VMware tags	NetBackup can include or exclude virtual machines based on the user assigned tags.

Support and use of VMware tag associations

NetBackup supports using VMware tags for virtual machine selection. You use this feature when you configure VMware Intelligent Policies to protect virtual machines. More information about this feature is available.

See “Notes and limitations for tag usage in VMware Intelligent Policy queries” on page 46.

See “Query rules for tags” on page 140.

NetBackup also supports the backup and restore of VMware tag associations with virtual machines when you use VMware Intelligent Policies. The tag association metadata for all tags that are associated with a virtual machine are backed up with that virtual machine. If those tags exist on the vCenter Server, they are recreated when the virtual machine is restored. More information about this feature is available.

See “Notes and limitations for the backup and restore of VMware tag associations” on page 47.

See “How NetBackup handles VMware tag associations at restore” on page 308.

The basics of a NetBackup query rule

For automatic virtual machine selection, NetBackup uses query rules to determine which VMware virtual machines to select for backup. You create the rules in the Query Builder, on the **Clients** tab of the policy.

Note: In the NetBackup web UI, you must use OData keywords and OData operators in query rules:

See "Query builder field reference" on page 141.

A query rule consists of the following:

- A keyword, such as **Displayname** or **Datacenter** (many keywords are available). For example: For automatic selection of the virtual machines with the display names that contain certain characters, you need the **Displayname** keyword in the rule.
- An operator, such as **Contains**, **StartsWith**, or **Equal**.
  The operator describes how NetBackup analyzes the keyword. For example: **Displayname StartsWith** tells NetBackup to look for the display names that start with particular characters.
- Values for the keyword.
  For the **Displayname** keyword, a value might be "prod". In that case, NetBackup looks for the virtual machines that have the display names that include the characters prod.
- An optional joining element (AND, AND NOT, OR, OR NOT) to refine or expand the query.

The policy uses these elements to discover and select virtual machines for backup.

Table 8-2 contains the examples of rules.

Table 8-2 Examples of rules

Rule	OData rule *	Description
`Displayname Contains "vm"`	`contains(displayName, 'vm')`	NetBackup selects the virtual machines that have the characters `vm` anywhere in their display names.
`Displayname EndsWith "vm"`	`endswith(displayName, 'vm')`	NetBackup selects the virtual machines that have the characters `vm` at the end of their display names.
`Datacenter AnyOf "datacenter1","datacenter2"`	`datacenter in ('datacenter1','datacenter2')`	NetBackup selects the virtual machines that use datacenter1 or datacenter2.

Table 8-2 Examples of rules *(continued)*

Rule	OData rule *	Description
`Powerstate Equal poweredOn`	`powerState eq 'poweredOn'`	NetBackup selects only the virtual machines that are currently turned on.
`Powerstate Equal poweredOn AND Tag Equal "Production"`	`powerState eq 'poweredOn' and tagName eq 'Production'`	NetBackup selects only virtual machines that are currently powered on with the "Production" tag.

* Use OData keywords only when you build queries with the NetBackup web UI.

Important notes on automatic virtual machine selection

The Virtual Machine Intelligent Policy feature in NetBackup is a different approach to VMware virtual machine selection in the policy. It represents a paradigm shift in the way you select virtual machines for backup. As with all major changes, the effective use of this feature requires forethought, preparation, and care.

Table 8-3 Important notes on automatic virtual machine selection!

Note!	Explanation
Create rules carefully.... 	Instead of manually selecting virtual machines for backup, you create guidelines for automatic selection of virtual machines. The guidelines are called rules; you enter the rules in the policy's Query Builder. You make the rules, and NetBackup follows them. If the rules state: Back up all virtual machines with a host name that contains "prod", NetBackup does that. Any virtual machine that is added to the environment with a host name containing "prod" is automatically selected and backed up when the policy runs. Virtual machines with the names that do not contain "prod" are not backed up. To have other virtual machines automatically backed up, you must change the query rules (or create additional policies).
Changes to the virtual environment can affect backup times. 	If many virtual machines are temporarily added to your environment and happen to fall within the scope of the query rules, they are backed up. The backups can therefore run much longer than expected.

Table 8-3 Important notes on automatic virtual machine selection! *(continued)*

Note!	Explanation
Test the query rules. 	Test the query rules ahead of time. The policy includes a Test Query function for that purpose. It's important to verify that your query operates as expected. Otherwise, the query may inadvertently select too many or too few virtual machines. As an alternative, you can use the `nbdiscover` command to test a query. Refer to the *NetBackup Commands Reference Guide*. Note also: The policy's **Primary VM identifier** parameter can affect the automatic selection process. See "Effect of Primary VM identifier parameter on Selection column in Test Query results" on page 154.
A query test does not create the backup list. NetBackup creates the backup list when the backup runs. 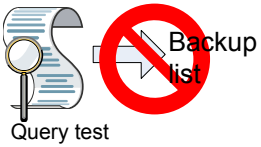 Query test  Backup execution	The automatic selection process is dynamic. Changes in the virtual environment may affect which virtual machines the query rules choose when the backup runs. **Note:** If virtual machine changes occur, the virtual machines that are selected for backup may not be identical to those listed in your query test results.
The policy does not display a list of the virtual machines that are to be backed up. Use the Activity monitor 	If you select virtual machines manually (with the Browse for Virtual machines screen), the selected virtual machines are listed on the policy **Clients** tab. But when you use the Query Builder for automatic selection, the selected virtual machines are not listed on the **Clients** tab. For a list of the backed up virtual machines, use the NetBackup Activity Monitor. See "Using the Activity monitor to monitor virtual machine backups" on page 222.

Table 8-3 Important notes on automatic virtual machine selection! *(continued)*

Note!	Explanation
When you save the policy, the query rules are not validated. 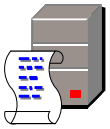 Policy	When you save a policy, policy validation does not consult the query rules and select virtual machines for backup. Because of the potential for changes in the virtual environment, virtual machine selection must wait until the backup runs. As a result, when you save the policy, NetBackup does not check the policy attributes against a backup list. If the query rules select the virtual machines that are incompatible with a policy attribute, policy validation cannot flag that fact. The incompatibility becomes apparent when NetBackup determines the backup list at the time of the backup. Take for example a policy that is configured for **Enable block-level incremental backup** (BLIB). BLIB works only with ESX 4.0 virtual machines at version vmx-07 or later. If the query rules select a virtual machine at a version earlier than vmx-07, the policy cannot back up that virtual machine. The mismatch between the policy and the virtual machine is revealed when the backup runs, not when the policy is validated. The Activity Monitor's job details log indicates which virtual machines can or cannot be backed up.

NetBackup requirements for automatic virtual machine selection

Note the following requirements for automatic selection of VMware virtual machines:

- The system where NetBackup runs must have access to the vCenter server.
- Automatic virtual machine selection requires no additional license beyond the NetBackup Enterprise Client license.
- Automatic virtual machine selection is required for backups in vCloud Director or for Replication Director for VMware.

http://www.veritas.com/docs/000033647

Automatic virtual machine selection: Task overview

This topic is a high-level overview of how to set up a NetBackup policy for automatic selection of VMware virtual machines. Follow the links in the table for more details.

Table 8-4 Automatic selection of virtual machines: overview of the tasks

Steps to configure automatic selection	Description and notes
Configure a VMware policy.	Use the policy **Attributes** tab. See "Configure a VMware policy" on page 87.
Set rules for virtual machine selection in the policy Query builder.	On the policy **Clients** tab, select **Select automatically through VMware intelligent policy query**. Choose a host for virtual machine selection (the default is the VMware backup host). To add rules, use the Query builder fields. See "Configure automatic virtual machine selection" on page 125. See "Options for selecting VMware virtual machines" on page 121.
Test the rules.	On the **Clients** tab, select **Test query**. Virtual machines are labeled as included or excluded, based on the rules. **Note:** The list of virtual machines is not saved in the **Clients** tab.
Run a backup.	When the policy runs, NetBackup consults the rules in the Query builder, creates a list of virtual machines, and backs them up.
Monitor the backup.	To see which virtual machines were backed up, use the Activity monitor. See "Using the Activity monitor to monitor virtual machine backups" on page 222.

Options for selecting VMware virtual machines

This topic describes the options on the policy **Clients** tab.

You can use these options to manually select virtual machines, or to configure NetBackup to select virtual machines automatically. For automatic selection, you specify the selection criteria (rules) in the policy's Query Builder. When the backup job runs, NetBackup discovers the virtual machines that currently meet the criteria and backs up those virtual machines.

See "Configure automatic virtual machine selection" on page 125.

Table 8-5 Virtual machine selection

Option	Description
Select manually	Click this option and click **Add** to manually enter virtual machines names, or to browse and select them from a list. See "Browse for VMware virtual machines" on page 107. **Note:** The rest of the fields and options are for automatic selection of virtual machines.
Select automatically through VMware intelligent policy query	Click this option to allow NetBackup to automatically select virtual machines for backup based on the rules that you enter in the Query Builder.
Enable VMware Cloud Director integration	Enables the backup of the virtual machines that reside in a vCloud environment. This option requires the automatic selection of virtual machines. With this option, the policy selects for backup only the virtual machines that vCloud manages: it skips the virtual machines that are not in vCloud.
NetBackup host to perform automatic virtual machine selection	This host discovers virtual machines and automatically selects them for backup based on your query rules. The resulting list determines which virtual machines are backed up. To designate your media servers as discovery hosts, select **Backup media server** from the pull-down.

Table 8-6 Query Builder

Option	Description
Basic mode	Places the Query Builder in Basic mode. See "Query builder field reference" on page 141.
Advanced mode	Places the Query Builder in Advanced mode for manual entry of rules. See "Using the Query builder in Advanced mode" on page 128. See "Query builder field reference" on page 141. See "Examples for the NetBackup Query Builder" on page 130.
Query Builder (Join, Field, Operator, Value)	Click **Add query**. Use the lists to select the values and define the rules for automatic selection of virtual machines. Click **Add** to add the list of queries. See "Query builder field reference" on page 141. See "Examples for the NetBackup Query Builder" on page 130.

Table 8-6 Query Builder *(continued)*

Option	Description
Actions > Edit	Use this option to change an existing query rule when in Basic mode, as follows: ■ Click the rule and then click **Edit**. ■ Make new selections in the Query Builder pull-down fields. ■ Click **Save**.
Actions > Remove	Removes a query rule when in Basic mode.
Test query	Click this option to test which virtual machines NetBackup selects based on the rules in the Query Builder. **Note:** This test option does not create the backup list for the policy. When the next backup runs from this policy, NetBackup rediscovers virtual machines and consults the query rules. At that time, NetBackup backs up the virtual machines that match the rules. See "Test Query screen for VMware" on page 152.
Reuse VM selection query results for	Sets the refresh rate of an internal cache of the query results. NetBackup uses the cache to determine which virtual machines to select at the time of the backup. The cache speeds up the selection of virtual machines without burdening the vCenter server at each scheduled backup. A faster cache refresh rate synchronizes the cache with the changes recorded in vCenter, such as the addition or removal of virtual machines. However, each cache refresh consumes vCenter resources. With a slower refresh rate, new virtual machines may not be included immediately in the next backup. New or changed virtual machines are included when the cache is refreshed. Note that fewer vCenter resources are consumed with a slower refresh rate. The default is 8 hours. For 8 hours, NetBackup uses the cache and does not attempt to rediscover virtual machines. Changes to the virtual environment do not affect the cache during that period. After 8 hours, when the policy runs next NetBackup rediscovers the virtual machines. If any changes match a rule in the query, the list of selected virtual machines is modified accordingly. **Note:** The cache is refreshed before the next scheduled backup whenever the policy is changed and saved. More information is available on the **Reuse VM selection query results for:** option: See "About the Reuse VM selection query results option" on page 124.

About the Reuse VM selection query results option

The NetBackup Test Query screen lists the virtual machines that NetBackup discovered in your virtual environment. Because the automatic selection feature is dynamic, later changes in the environment may affect which virtual machines match the query rules. For example: if virtual machines are added later, the current test results may not be identical to the virtual machines that are selected when the backup runs.

During the period you specify on **Reuse VM selection query results for**, NetBackup reuses the current list of virtual machines as the backup list. It does not consult the Query Builder or rediscover virtual machines.

The less often your virtual machine environment undergoes changes, the more advantageous it may be to reuse the list of virtual machines for backups. In large environments, discovery of virtual machines takes time and consumes resources on the vCenter server.

Note the following about the Reuse VM selection query results option:

- Determines how long the query results are reused (that is, how often the list of discovered virtual machines is refreshed).
- Controls how often NetBackup performs discovery on the vCenter server.
  For the environments that contain many virtual machines, the discovery process may increase the load on the vCenter server.
- Has no effect on when the NetBackup policy schedule runs.
- Is invalidated if the query rules are changed or if the policy attributes that affect discovery are changed. In that case, NetBackup rediscovers virtual machines the next time the policy runs.
  The following topic describes the policy attributes that affect discovery:
  See “Effect of Primary VM identifier parameter on Selection column in Test Query results” on page 154.

For example: assume that the **Reuse VM selection query results for** option is set to 8 hours and your query selects turned-on virtual machines. If additional virtual machines are turned on during the 8-hour period, they are not added to the policy's backup list. The policy backs up the virtual machines that were last added to the list (such as when the policy was created). After 8 hours, when the policy runs next, the recently turned on virtual machines are discovered and added to the backup list.

If the next backup occurs before the Reuse period expires, and a virtual machine was renamed during the Reuse period, NetBackup backs up the renamed virtual

machine under its original name. Because the reuse period has not expired, NetBackup does not rediscover virtual machines and therefore cannot identify the virtual machine by its new name. (NetBackup identifies the virtual machine by its instance UUID.) To have backed up the virtual machine with its new name, the Reuse period should have been set to a shorter interval.

Note: The virtual machines that have been selected for backup are not displayed on the policy **Clients** tab. To see which virtual machines NetBackup has selected and backed up, refer to the following topics.

See "Using the Activity monitor to monitor virtual machine backups" on page 222.

The effect of virtual machine discovery on vCenter

In all but the largest environments, it may be advantageous to set the **Reuse VM selection query results for** option so that discovery occurs more often.

If changes occur to virtual machine configuration (such as adding, deleting, or moving vmdk files), it may be necessary to set **Reuse VM selection query results for** to 0. With a setting of 0, NetBackup rediscovers the virtual machines and their configuration each time the policy runs.

Note: If the vmdk files are reconfigured and the next backup runs without rediscovery, NetBackup is not aware of the vmdk changes. It attempts to back up the virtual machines in their previous configuration. The result may be an incorrect backup.

Configure automatic virtual machine selection

NetBackup can automatically select VMware virtual machines for backup based on the criteria that you enter. You specify the criteria (rules) in the Query Builder on the NetBackup policy **Clients** tab. You can set up rules to include certain virtual machines for backup, or to exclude virtual machines.

Note: In the NetBackup web UI, you must use OData keywords and OData operators in query rules:

See "Query builder field reference" on page 141.

When the backup job runs, NetBackup creates a list of the virtual machines that currently meet the query rules and backs them up.

The Query Builder can operate in Basic mode or in Advanced mode.

To configure automatic virtual machine selection in Basic mode

1. On the policy **Attributes** tab, select **VMware** for the policy type.
2. On the policy **VMware** tab, select a VMware backup host.

   Review the other options on the **VMware** tab.

   See "Backup options on the VMware tab" on page 90.
3. Make other policy selections as needed (for example, create a schedule).
4. Select the **Clients** tab, and select **Select automatically through VMware intelligent policy query**.

   If you selected virtual machines manually, those virtual machines are removed from the policy.
5. To back up virtual machines in vCloud Director, select **Enable VMware Cloud Director integration**.

   Note: Enable VMware Cloud Director integration makes several vCloud Director keywords available in the policy Query Builder Field, for rule-based selection of virtual machines. If this option is not selected, NetBackup cannot use the vCloud keywords to locate virtual machines in vCloud Director, and the backup fails.

   ---
6. To create a rule, make selections from the menus.

   For the first rule, you can start with the **Field** list, depending on the type of rule. (For the first rule, the only selections available for the **Join** field are blank (none), or NOT.)

   Then make a selection for **Operator**.

   For the **Value** field: Click the folder icon to browse for values, enter the value manually, or in some cases use the **Value** drop-down. The characters you enter manually in the **Value** field must be enclosed in single quotes or double quotes. Note that browsing for values may take some time in large virtual environments.

   See "Query builder field reference" on page 141.
7. Click the **Add** to add the rule to the **Query** pane.
8. Create more rules as needed.

   See "Query builder field reference" on page 141.

   See "Examples for the NetBackup Query Builder" on page 130.

9 To see which virtual machines NetBackup currently selects based on your query, click **Test query**.

The Virtual machines in your current environment that match the rules for selection in the policy are labeled INCLUDED. Note however that the Test query option does not create the backup list for the policy. When the next backup runs from this policy, NetBackup rediscovers virtual machines and consults the query rules. At that time, NetBackup backs up the virtual machines that match the query rules.

The list of virtual machines is saved but the virtual machines are not displayed in the policy's **Clients** tab.

See "Test Query screen for VMware" on page 152.

10 You can specify how long NetBackup uses the latest query results as the backup list for future executions of the policy. Set the time period in **Reuse VM selection query results for**.

11 To create queries manually instead of using the menus, click **Advanced mode**.

See "Using the Query builder in Advanced mode" on page 128.

Editing an existing query in Basic mode

You can use the Query Builder to enter rules for the automatic selection of VMware virtual machines for backup.

Note: In the NetBackup web UI, you must use OData keywords and OData operators in query rules:

See "Query builder field reference" on page 141.

To edit an existing query in Basic mode

1 Locate on the query rule you want to change and click **Actions > Edit**.

2 Make selections in the menus.

3 Click **Save**.

4 To delete a rule, locate the rule and click **Actions > Remove**.

See "Using the Query builder in Advanced mode" on page 128.

Using the Query builder in Advanced mode

You can use the Query builder to enter rules for the automatic selection of VMware virtual machines for backup.

The Query builder's Advanced mode provides more flexibility in crafting rules for virtual machine selection, including the use of parentheses for grouping.

Note: In the NetBackup web UI, you must use OData keywords and OData operators in query rules:

See "Query builder field reference" on page 141.

To use the Query builder in Advanced mode

1. Set up a VMware policy and specify a VMware backup host or backup media server.

   For assistance, you can refer to the first few steps of the following procedure:

   See "Configure automatic virtual machine selection" on page 125.

2. Click the **Clients** tab.
3. Click **Select automatically through VMware intelligent policy query**.
4. Locate **Query builder** and select **Advanced mode**.
5. You can use the Query builder menus to add query rules. You can also type in rules manually.

   Here are a few example queries:

   ```
   VMFolder Contains "mango"
   ```

   ```
   Datastore StartsWith "Acc" OR Datastore StartsWith "Prod"
   ```

   ```
   vCenter Contains "ROS" AND ESXserver Equal "VM_test1" AND
   Powerstate Equal poweredOn
   ```

6 To insert a rule between existing rules, place the cursor where you want the new rule to start and type it in.

When you create a rule with the drown-down menus, it appears at the end of the query. You can cut and paste it into the proper location.

7 To establish the proper order of evaluation in compound queries, use parentheses to group rules as needed. Compound queries contain two or more rules, joined by AND, AND NOT, OR, or OR NOT.

See "AND vs. OR in queries" on page 129.

See "Order of operations in queries (precedence rules)" on page 135.

See "Parentheses in compound queries" on page 136.

AND vs. OR in queries

The **Join** field in the Query Builder provides connectors for joining rules (AND, AND NOT, OR, OR NOT). The effect of AND versus OR in the Query Builder may not be obvious at first glance.

In essence, AND and OR work in this way:

- AND limits or restricts the scope of the query.
- OR opens up the query to an additional possibility, expanding the scope of the query.

Note: Do not use AND to join the rules that are intended to include additional virtual machines in the backup list. For instance, AND cannot be used to mean "include virtual machine X AND virtual machine Y."

For example: To include the virtual machines that have either "vm1" or "vm2" in their names, use OR to join the rules:

```
Displayname Contains "vm1"
OR Displayname Contains "vm2"
```

If you use AND to join these rules:

```
Displayname Contains "vm1"
AND Displayname Contains "vm2"
```

the result is different: the backup list includes only the virtual machines that have both vm1 and vm2 in their names (such as "acmevm1vm2"). A virtual machine with the name "acmevm1" is not included in the backup.

Table 8-7 provides the examples with AND and OR.

Table 8-7 Queries with AND, OR

Query	Description
Displayname Contains "vm1" OR Displayname Contains "vm2"	This query selects any virtual machine that has either vm1 or vm2 in its display name. For example, this query selects both "seabizvm1" and "seabizvm2" for backup.
vCenter Equal "vCenterServer_1" AND Datacenter Equal "dc_A" AND ESXserver Equal "prod" AND VMHostName Contains "manu"	This query is very specific. Virtual machines with the host names that contain "manu" are included in the backup only if: they reside in vCenter server "vCenterServer_1", datacenter "dc_A", and ESX server "prod". The virtual machines that do not reside in that hierarchy are not included. For example: if a virtual machine resides in "vCenterServer_1" and datacenter "dc_A", but not in ESX server "prod", that virtual machine is not included.
vCenter Equal "vCenterServer_1" OR Datacenter Equal "dc_A" OR ESXserver Equal "prod" OR VMHostName Contains "manu"	This query uses the same keywords and values, but combines them with OR. The result may be a much larger list of virtual machines. A virtual machine that meets any of these rules is included: ■ Any virtual machines in vCenter "vCenterServer_1". Their host names, datacenter, or ESX server do not matter. ■ Any virtual machines in datacenter "dc_A". Their host names or server do not matter. ■ Any virtual machines in ESXserver "prod". Their host names, datacenter, or vCenter server do not matter. ■ Any virtual machines with a host name that contains "manu". Their server or datacenter do not matter.

Examples for the NetBackup Query Builder

The following table provides example query rules.

To use the Query Builder, you must click **Select automatically through VMware intelligent policy query** on the **Clients** tab.

Click **Advanced mode** to see the query rule in Advanced mode. Only this mode supports the use of parentheses for grouping sets of rules.

See "Using the Query builder in Advanced mode" on page 128.

Note: The advanced mode of the Query Builder uses OData keywords and operators.

See "Query builder field reference" on page 141.

Another topic is available on the difference between AND and OR in a query.

See “AND vs. OR in queries” on page 129.

Table 8-8 Query Builder examples

VIP Example query	OData example query *	Query result when backup job runs
No query rules specified (Query pane is empty)	No query rules specified (Query pane is empty)	All virtual machines are added to the backup list. Exceptions are those that do not have a host name, or that have invalid characters in the display name. See “Effect of Primary VM identifier parameter on Selection column in Test Query results” on page 154.
Displayname Contains "prod"	contains(displayName, 'prod')	All virtual machines with the display names that contain the string "prod" are added to the backup list. See “Effect of Primary VM identifier parameter on Selection column in Test Query results” on page 154.
powerstate Equal "poweredOn"	powerState eq 'poweredOn'	Any virtual machine that is turned on is added to the backup list.
VMGuestOS Equal "windows7Guest"	vmGuestOs eq 'windows7Guest'	All virtual machines with a guest OS of Windows 7 are added to the backup list.
DisplayName AnyOf "grayfox7”,"grayfox9"	displayName in ('grayfox7', 'grayfox9')	The virtual machines named "grayfox7”and "grayfox9" are added to the backup list. (Note that each value must be enclosed in its own quotes, with a comma in between.)
powerstate Equal "poweredOn" AND Datastore Equal "Storage_1" AND VMGuestOS Equal "rhel4Guest"	powerState eq 'poweredOn' and datastoreName eq 'Storage_1' and vmGuestOs eq 'rhel4Guest'	In datastore Storage_1: any virtual machine that is turned on and has a guest OS of Red Hat Linux 4 is added to the backup list.
vCenter Equal "vCenterServer_1" AND ESXserver Contains "prod"	vCenter eq 'vCenterServer_1' and contains(host, 'prod')	In the vCenter server vCenterServer_1, virtual machines that are in ESX servers with names containing "prod" are added to the backup list.
Cluster Equal "VMcluster_1" AND ESXserver AnyOf "ESX_1","ESX_2","ESX_3" AND VMHostName Contains "Finance"	cluster eq 'VMcluster_1' and host in ('ESX_1', 'ESX_2', 'ESX_3') and contains(hostName, 'Finance')	In cluster VMcluster_1, all virtual machines with the host names that contain "Finance", in ESX servers ESX_1, ESX_2, ESX_3, are added to the backup list.

Table 8-8 Query Builder examples *(continued)*

VIP Example query	OData example query *	Query result when backup job runs
VMFolder StartsWith "Prod" OR VMFolder NotEqual "VM_test"	startswith(vmFolder, 'Prod') or vmFolder ne 'VM_test'	For any folder whose name starts with "Prod" or whose name is not "VM_test", add its virtual machines to the backup list.
Examples with IsSet		See "The IsSet operator in queries" on page 132.
Datacenter Contains "prod" AND Tag Equal "Finance"	contains(datacenter, 'prod') and tagName eq 'Finance'	This query selects any virtual machine where the Datacenter contains "prod" and the user-specified tag is "Finance".
Datacenter Equal "prod" AND NOT Tag Equal "Test"	datacenter eq 'prod' and not (tagName eq 'test')	Selects virtual machines where Datacenter is "prod" but excludes any virtual machines that have the user-specified tag "Test".

* Use OData operators only when you build queries with the NetBackup web UI's advanced mode under the Query Builder or with NetBackup APIs.

The IsSet operator in queries

In a query, you can use the IsSet operator to ensure that certain virtual machines are included or excluded from the backup.

For example: if the **Primary VM identifier** parameter is set to VM hostname, NetBackup is unable to identify virtual machines for backup that do not have a host name. You can use IsSet to exclude such virtual machines from the backup list.

Table 8-9 Examples of queries with the IsSet operator

Query rules with IsSet operator	OData query rules with IsSet operator *	Effect of the query on virtual machine selection
Cluster Contains "dev" AND VMDNSName IsSet	contains(cluster, 'dev') and dnsName ne null	INCLUDED: Any virtual machine in a cluster that has a name that contains the string "dev" if the virtual machine also has a VMware DNS name. EXCLUDED: Any virtual machines that do not have a VMware DNS Name. Without VMDNSName IsSet in this query, virtual machines without a DNS name cannot be excluded. They would be listed as FAILED.
Displayname Contains "prod" AND VMHostName IsSet	contains(displayName, 'prod') and hostName ne null	INCLUDED: Any virtual machine with a display name that contains the string "prod" if the virtual machine also has a host name. EXCLUDED: Any virtual machines that do not have host names. Without VMHostName IsSet in this query, virtual machines without a host name cannot be excluded. They would be listed as FAILED.

* Use OData operators only when you build queries with the NetBackup web UI.

The policy's **Primary VM identifier** parameter has an important effect on which virtual machines NetBackup can back up. This parameter affects the test query results.

See "Effect of Primary VM identifier parameter on Selection column in Test Query results" on page 154.

About selecting virtual machines by means of multiple policies

If your virtual environment has many virtual machines with inconsistent naming conventions, you may need multiple policies working in tandem. It may be difficult to create a single policy that automatically selects all the virtual machines that you want to back up.

For this situation, configure several policies such that each policy backs up a portion of the environment. One policy backs up a particular set or group of virtual machines, such as those that have host names. A second policy backs up a different group of virtual machines that were not backed up by the first policy, and so forth. When all the policies have run, all the virtual machines are backed up.

The following table describes the policies that are designed to back up the virtual environment in three phases. Note that each policy relies on a different setting for the **Primary VM identifier** parameter.

Table 8-10 Three policies that back up virtual machines in phases

Policy	Query Builder rules	OData Query Builder rules *	Backup result
First policy **Primary VM identifier** parameter: VM hostname	VMHostName IsSet	hostName ne null	This policy backs up all virtual machines that have a host name. Any virtual machines that do not have a host name are excluded from the backup.
Second policy **Primary VM identifier** parameter: VM display name	NOT VMHostName IsSet AND VMHasVDSName Equal 'TRUE'	not (hostName ne null) and vmHasVdsName eq 'TRUE'	This policy backs up all virtual machines that do not have a host name but that do have a valid display name. Any virtual machines that do not have a host name or a valid display name are excluded from the backup. See "NetBackup character restrictions for the Primary VM identifier" on page 42.
Third policy **Primary VM identifier** parameter: VM UUID	NOT VMHostName IsSet AND NOT VMHasVDSName Equal 'TRUE'	not (hostName ne null) and not(vmHasVdsName eq 'TRUE')	This policy backs up the virtual machines that were not backed up by the first two policies. This policy selects the virtual machines that do not have a host name or a valid display name, but that do have a UUID.

* Use OData operators only when you build queries with the NetBackup web UI.

More information is available on the **Primary VM identifier** parameter and its effect on virtual machine selection.

See "Effect of Primary VM identifier parameter on Selection column in Test Query results" on page 154.

See "The basics of a NetBackup query rule" on page 117.

Order of operations in queries (precedence rules)

The information in this topic is for advanced users who understand precedence in programming languages. In the Query Builder, the order in which operations occur can determine which virtual machines are selected and backed up.

The following table lists the order of operations, or precedence, from highest to lowest (7 is the highest). For example, an operation with a precedence of 6 (such as Contains) is evaluated before an operation with a precedence of 5 (such as Greater).

Table 8-11 Order of operations

Operation	Description	Precedence
!x	Produces the value 0 if x is true (nonzero) and the value 1 if x is false (0).	7
x Contains y	Does y exist somewhere in x	6
x StartsWith y	Does x start with y	6
x EndsWith y	Does x end with y	6
x AnyOf list	Does x appear in list	6
x Greater y	Is x greater than y	5
x GreaterEqual y	Is x greater than or equal to y	5
x Less y	Is x less than y	5
x LessEqual y	Is x less than or equal to y	5
x Equal y	Is x equal to y	4
x NotEqual y	Is x not equal to y	4
Not x	operator produces the value 0 if x is true (nonzero) and the value 1 if x is false (0).	3
x And y	True if both x and y are true	2
x OR y	True if either x or y are true	1

Note the following:

- AND has a higher precedence than OR.
  In the Query Builder's Advanced Mode, you can use parentheses to change the order of evaluation in the rules that use AND or OR.

See “Parentheses in compound queries” on page 136.

- In the Query Builder's Advanced Mode, you can combine two or more operations in a single rule without AND or OR to join them. Precedence determines the order in which the operations are evaluated within the rule.
  Example of a rule that includes three operations:

  ```
  Displayname StartsWith “L” NotEqual Displayname contains “x”
  ```

  This rule selects the following virtual machines:
  Virtual machines with the names that start with L but do not contain x.
  Virtual machines with the names that do not start with L but that do contain x.
  Explanation: The StartsWith and Contains operations have a precedence of 6, whereas NotEqual has a lower precedence of 3. Starting on the left, the StartsWith operation is evaluated first and the Contains operation is evaluated next. The last operation to be evaluated is Not Equal.

See “Using the Query builder in Advanced mode” on page 128.

Parentheses in compound queries

You can use the Query Builder to make precise queries containing as many rules as necessary to identify the appropriate virtual machines. In a query such as powerstate Equal "poweredOn", the result of the query is easy to predict: only the virtual machines that are turned on are included in the backup. But if several rules are combined with AND and OR, the result may not be obvious. This kind of query is called a compound query. Compound queries contain two or more rules, joined by AND, AND NOT, OR, or OR NOT.

The order in which the Query Builder evaluates compound rules affects the outcome of the query. Grouping the rules with parentheses can change the order of evaluation and thus the outcome of the query.

The examples in the following table demonstrate how the Query Builder evaluates compound queries with and without parentheses.

Note: Only the Query Builder's Advanced Mode supports the use of parentheses.

Table 8-12 Examples of compound queries with and without parentheses

Example query	The following virtual machines are selected
ESXServer Equal "ESX001" OR Folder Equal "FolderEngA" AND powerstate Equal ON	All virtual machines under ESX001 (regardless of power state), and virtual machines under FolderEngA that are turned on To select only the virtual machines that are turned on in the ESX server and in the folder, use parentheses (see next example).
(ESXServer Equal "ESX001" OR Folder Equal "FolderEngA") AND powerstate Equal ON	All the virtual machines that are turned on in ESX001 and in FolderEngA.

Query rules for resource pools

If the resource pool that you query is nested, the choice of Operator determines which virtual machines in the resource pool hierarchy are discovered.

For example, assume the following hierarchy of resource pools that contain virtual machines:

```
Res/ResourcePool_1
                VM1
                VM2
                /ResourcePool_2
                            VM3
                            VM4
                            /ResourcePool_3
                                        VM5
                                        VM6
```

where ResourcePool_1 contains virtual machines VM1 and VM2, and so forth.

The following table shows the query results with the Contains, Equal, StartsWith, and EndsWith operators. (Other operators can be used.)

Note: If you want the query to include all virtual machines in a hierarchy of nested resource pools, do not use Equal as the Operator.

Table 8-13 Example rules for nested resource pools

Query rule	Included virtual machines
Resourcepool Contains "Res/ResourcePool_1"	Includes all the virtual machines in the three resource pools (VM1 through VM6).

Table 8-13 Example rules for nested resource pools *(continued)*

Query rule	Included virtual machines
Resourcepool Equal "Res/ResourcePool_1"	Includes only the virtual machines that are in ResourcePool_1 (VM1, VM2). Virtual machines in the sub-pools are not included (VM3 through VM6).
Resourcepool Equal "Res/ResourcePool_1/ResourcePool_2"	Includes only the virtual machines that are in ResourcePool_2 (VM3, VM4).
Resourcepool StartsWith "Res/ResourcePool"	Includes all the virtual machines in the three resource pools (VM1 through VM6).
Resourcepool StartsWith "Res/ResourcePool_1/ResourcePool_2"	Includes only the virtual machines that are in ResourcePool_2 and 3. Virtual machines in ResourcePool_1 are not included.
Resourcepool EndsWith "ResourcePool_2"	Includes the virtual machines in ResourcePool_2 (VM3, VM4) but not in ResourcePool_1 or 3.

These examples also apply to host folders.

See “Query rules for datacenter folders (host folder)” on page 138.

Query rules for datacenter folders (host folder)

In NetBackup terminology, a host folder is a folder that has been defined within a VMware datacenter. A host folder can contain ESX servers or clusters, as well as other folders. For example:

```
Folder_1
       ESX1
       ESX2
       subfolder_A
                 ESX3
                 ESX4
       subfolder_B
                 ESX_5
                 ESX_6
```

If you want NetBackup to select all the virtual machines within the top-level folder and any subfolders, use the Contains or StartsWith operator. For example:

HostFolder Contains "Folder_1"

Note: If you want the query to include all virtual machines in the hierarchy of folders, do not use Equal as the Operator.

If you want NetBackup to select the virtual machines within a subfolder only (such as subfolder_A), use the Contains or Equal operator. For example:

HostFolder Equal "Folder_1/subfolder_A"

In this case, NetBackup includes only the virtual machines that reside on servers ESX3 and ESX4.

For host folders, these operators work the same as they do for resource pools. For further query builder examples, refer to the following topic (substitute host folder for resource pool in the examples):

See "Query rules for resource pools" on page 137.

Note: If an ESX cluster is not contained within a folder and you click the browse button: The ESX cluster name appears in the **List of possible values for *Value*** dialog.

Query rules for duplicate names

If you have clusters, datastores, or virtual machine display names that have duplicates elsewhere in your virtual environment, note: The query rules must specify the parent datacenter or host folder to avoid conflicts during discovery. (A host folder is one that has been defined within a datacenter.)

Take the following example of duplicate virtual machine names:

```
Folder_1
       ESXi_prod
              VM_1
Folder_2
       ESXi_mrkt
               VM_1
```

To back up `ESXi_prod/VM_1` but not `ESXi_mrkt/VM_1`, use the following query:

```
HostFolder Equal "Folder_1"
AND ESXserver Equal "ESXi_prod"
AND Displayname Equal "VM_1"
```

To back up only `ESXi_mrkt/VM_1`, use the following query:

```
HostFolder Equal "Folder_2"
AND ESXserver Equal "ESXi_mrkt"
AND Displayname Equal "VM_1"
```

Note: To back up both of these virtual machines from the same policy, include both of these rules in the policy, as follows:

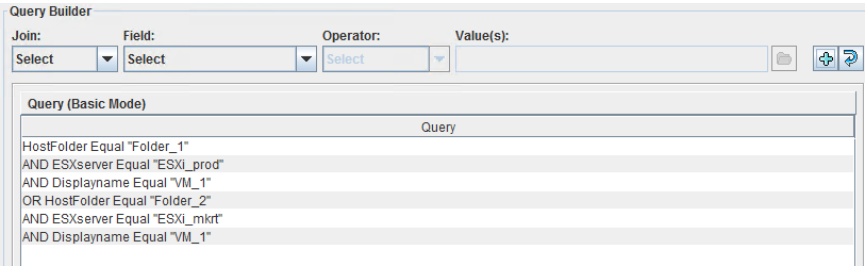

Instead of `Equal`, you can use other field values, such as `Contains`.

See "Query rules for datacenter folders (host folder)" on page 138.

Query rules for tags

Be aware of the following rules when you use **tags** in your queries:

- If you use VMware tags, you can base your backup selections on these tags. Be aware these tags are case-sensitive, so `UNIX` is different from `unix`.
- When NetBackup uses tags to select virtual machines, the selection is based only on tag names. The selection is independent of the category.
  Example:
  - `Virtual_machine_1` has a user-specified tag `HR` from the category `production`.
  - `Virtual_machine_2` has a user-specified tag `HR` from the category `test`.

  Queries that select virtual machines with the tag `HR` select both virtual machines.
- NetBackup uses a different VMware interface to access tag information from vCenter Servers than is used for other query fields. Therefore, tag related calls to the vCenter Server run only if they are required. Calls are skipped if the query is satisfied without the tag information. NetBackup only collects this information once per vCenter Server. NetBackup collects tag metadata as part of virtual machine backup, but it collects tag metadata only from those vCenter Servers for which a virtual machine is selected. If no virtual machine is selected from a vCenter Server and tags are not used in the query then tag metadata is not collected from that vCenter Server.

- Example 1: `Tag Equal "Production" OR Powerstate Equal poweredOn`
- Example 2: `Powerstate Equal poweredOn OR Tag Equal "Production"`

In Example 1, NetBackup retrieves virtual machine data as well as tag data from each virtual server for which it has credentials.

In Example 2, NetBackup retrieves virtual machine data for each virtual server for which it has credentials. But NetBackup only needs to retrieve tag data for virtual machines where the `Powerstate` is not equal to `poweredOn`.

Query builder field reference

You can use the Query builder to enter rules for the automatic selection of VMware virtual machines for backup.

Note: The advanced mode of the Query builder uses OData keywords and operators.

Table 8-14 describes the fields and options for creating rules in the Query builder.

Table 8-14 Query builder options: Join, Field, Operator, Value

Query Builder fields	Description
Join	Selects a connector to join the rules. For the first rule, the choices are blank (none) or NOT. After you add a rule, the available connectors are AND, AND NOT, OR, OR NOT.
Field	Selects a parameter on which to build the rule. Select one from the list (scroll down for additional parameters). You can type the first character to speed up selection. For example, on entering “d”, the list moves to first entry starting with "d". Another entry of "d" moves through the list to the next entry starting with "d". The selected entry is automatically filled in. See Table 8-15 on page 142.
Operator	Selects an operator. The available operators depend on the parameter that is selected for **Field**. See Table 8-16 on page 151.

Table 8-14 Query builder options: Join, Field, Operator, Value *(continued)*

Query Builder fields	Description
Value	Specifies a value for the **Field** parameter. The **Value** field allows manual entry. It may also be a dropdown, depending on the selections that are made in the other fields. For manual entry, you can specify multiple comma-separated values. See Table 8-17 on page 152.
Browse icon	Allows browsing for specific values, depending on the selections that are made in the other dropdown fields.
Save	Adds the current query selections to the **Query** pane as a new rule.

Field (keywords)

Table 8-15 describes the keywords available in the **Field** dropdown. The table also indicates whether the values for each keyword (in the **Value** field) are case-sensitive.

Note: Use OData **Field** keywords only when you build queries with the NetBackup web UI 's advanced mode under the Query Builder or with NetBackup APIs.

Keep in mind that the **Field** keyword alone does not determine the inclusion or exclusion of virtual machines. It depends on the rule that you construct: the combination of Join, Field, Operator, and Values.

Table 8-15 Keywords in the **Field** dropdown

Field keyword	OData field keyword	Data type	Description
Annotation	annotation	Alphanumeric string	The text that is added to virtual machine annotations in vSphere Client. Values are case-sensitive.
assetGroup	assetGroup		NEW

Table 8-15 Keywords in the **Field** dropdown *(continued)*

Field keyword	OData field keyword	Data type	Description
cluster	cluster	Alphanumeric string	The name of the cluster (a group of ESX servers) that the virtual machine is configured in. Values are not case-sensitive. **Note:** A virtual machine may be assigned to an ESX server that is not in a cluster. Note also that in VMware, a cluster name need only be unique within a datacenter path.
datacenter	datacenter	Alphanumeric string	The name of the VMware datacenter. Values are not case-sensitive.
datacenterPath	datacenterPath	Alphanumeric string	The folder structure that defines the path to a datacenter. Use this option if the datacenter name that you want to filter on is not unique in your environment. Values are case-sensitive.
datastore	datastoreName	Alphanumeric string	The name of the datastore. Values are case-sensitive. **Note:** Multiple ESX servers can share access to the same datastore. Also, a datastore name can be duplicated between multiple ESX servers. Use DatacenterPath or ESXserver to uniquely identify the datacenter.
datastoreCluster	datastoreCluster	Alphanumeric string	The name of the datastore cluster that contains the datastores. Values are not case-sensitive.
datastoreFolder	datastoreFolder	Alphanumeric string	The name of the folder that contains the datastores. Values are not case-sensitive.
datastoreNfsHost	datastoreNfsHost	Alphanumeric string	The name of the datastore's NFS host. Values are not case-sensitive.

Table 8-15 Keywords in the **Field** dropdown *(continued)*

Field keyword	OData field keyword	Data type	Description
datastoreNfsPath	datastoreNfsPath	Alphanumeric string	The folder structure that defines the path to an NFS datastore. Use this option if the NFS host name of the datastore that you want to filter on is not unique in your environment. Values are not case-sensitive.
datastoreType	datastoreType	Alphanumeric string	The type of the datastore. Values are NFS, NFS41, VMFS, vsan, and VVOL. Values are not case-sensitive.
displayname	displayName	Alphanumeric string	The virtual machine's display name. Values are case-sensitive.
host	host	Alphanumeric string	The name of the ESX server. Values are not case-sensitive. The ESX host name must match the name as defined in the vCenter server.
hostFolder	hostFolder	Alphanumeric string	The folder path between the datacenter level and a cluster, ESX hosts, or a subfolder. If an ESX cluster is not contained within a folder and you click the browse button: The ESX cluster name appears in the **List of possible values for** dialog. Values are not case-sensitive. See "Query rules for datacenter folders (host folder)" on page 138.
instanceUuid	instanceUuid	Alphanumeric string	The instance UUID of the virtual machine. Example query: `InstanceUUID Equal "501b13c3-52de-9a06-cd9a-ecb25aa975d1"` Values are not case-sensitive.
networkName	network	Alphanumeric string	The name of the network switch (on an ESX server) or distributed switch. Values are not case-sensitive.

Table 8-15 Keywords in the **Field** dropdown *(continued)*

Field keyword	OData field keyword	Data type	Description
networkFolder	networkFolder	Alphanumeric string	The name of the folder that contains the network. Values are not case-sensitive.
powerstate	powerState	Alphabetic	The power state of the virtual machine. Values are poweredOff, poweredOn, suspended.
resourcePool	resourcePool	Alphanumeric string	The name of the resource pool. (A resource pool is similar to a vApp.) Values are case-sensitive. If a resource pool contains other resource pools (sub-pools), the choice of Operator determines whether virtual machines in the sub-pools are included. See "Query rules for resource pools" on page 137.
tagName	tagName	Alphanumeric string	The name of the tag. Values are case-sensitive. When NetBackup uses tags to select virtual machines, the selection is based only on tag names. The selection is independent of the category. See "Query rules for tags" on page 140.
template	template	Boolean	TRUE if the virtual machine is a virtual machine template.

Table 8-15 Keywords in the **Field** dropdown *(continued)*

Field keyword	OData field keyword	Data type	Description
vApp	vApp	Alphanumeric string	The name of the vApp. Values are case-sensitive. A vApp is a collection of virtual machines. vApps can also contain resource pools and other vApps. vApps are components of standalone ESX servers or of clusters. Like vSphere Client, NetBackup refers only to the top level of a vApp that contains sub vApps. For the following rule: `vApp Equal "vapp1"` If vapp1 has a sub vApp named "vapp2", any virtual machines in vapp1 or vapp2 are included. You cannot make a rule that refers specifically to vapp2.
vCDCatalog	vcdCatalog	Alphanumeric string	The name of the vCloud Director catalog. Values are not case-sensitive.
vCDIsExpired	vcdIsExpired	Alphabetic	Expired if the vCloud Director vApp or vApp template is expired. Possible values are Expired, Not Expired, and Unknown. Unknown indicates that an error occurred between vCloud Director and the vSphere environment. Examples are errors in provisioning or in a deletion operation.
vCDIsvAppTemplate	vcdIsvAppTemplate	Boolean	TRUE if the vCloud Director vApp is a template.
vCDOrg	vcdOrg	Alphanumeric string	The name of the vCloud Director organization. Values are not case-sensitive.
vCDOrgvDC	vcdOrgvDC	Alphanumeric string	The name of the organization virtual datacenter in vCloud Director. Values are not case-sensitive.
vCDServer	vcdServer	Alphanumeric string	The name of the vCloud Director server. Values are not case-sensitive.

Table 8-15 Keywords in the **Field** dropdown *(continued)*

Field keyword	OData field keyword	Data type	Description
vCDvApp	vcdVapp	Alphanumeric string	The name of the vCloud Director vApp. Values are not case-sensitive.
vCenter	vCenter	Alphanumeric string	The name of the vCenter server. Values are not case-sensitive. The vCenter name that is specified in the Query Builder must match the name as entered when you added its credentials.) Note that a fully qualified domain name is recommended.
vCenterVersion	vCenterVersion	Alphanumeric string	The version of the vCenter Server. For example: 5.1.0, 5.5.0, 6.0.0 The possible values of this field are automatically updated and populated based on the environment. Only the versions of the vCenter servers that are registered with NetBackup are shown.
dnsName	dnsName	Alphanumeric string	The virtual machine DNS name in vSphere Client. Values are not case-sensitive.

Table 8-15 Keywords in the **Field** dropdown *(continued)*

Field keyword	OData field keyword	Data type	Description

vmFolder	vmFolder	Alphanumeric string	The name of the VM folder (within a datacenter), including the path to the folder that contains the VMs.
			Values are not case-sensitive.
			For example, assume the following VM folders containing a total of 65 VMs:
			`vm\VM_backup_prod1 (contains 5 VMs)`
			`vm\VM_backup_prod1\cluster1 (contains 10 VM`
			`vm\VM_backup_prod2 (contains 50 VMs)`
			To include the VMs in `vm\VM_backup_prod1` but not the VMs in `cluster1` or in any other folder:
			`VMFolder Equal "vm\VM_backup_prod1"`
			To include the VMs in `vm\VM_backup_prod1` and in its subfolder `cluster1`:
			`VMFolder Equal "vm\VM_backup_prod1" OR VMFolder StartsWith "vm\VM_backup_prod1\`
			Note: The first backslash is an escape character that causes the following backslash to be interpreted as a literal character.
			To include all 65 VMs:
			`VMFolder StartsWith "vm\VM_backup_prod"`
			Note: Any VM that is in a path that begins with `vm\VM_backup_prod` is included.

vmHasIde	vmHasIde	Boolean	TRUE if the virtual machine has IDE drives.
VMhasIndD	vmHasIndD		NEW
vmHasMds	vmHasMds	Boolean	TRUE if the virtual machine has multiple datastores. You can use this keyword to select any virtual machine that is configured to use more than one datastore.

Table 8-15 Keywords in the **Field** dropdown *(continued)*

Field keyword	OData field keyword	Data type	Description
VMHasNVME	vmHasNvme	Boolean	TRUE if the virtual machine has NVMe drives.
VMHasRDM	vmHasRdm	Boolean	TRUE if the virtual machine uses Raw Device Mapping (RDM).
VMHasRDMO	vmHasRdmo		NEW
VMHasSATA	vmHasSata	Boolean	TRUE if the virtual machine has SATA drives.
VMHasSnap	VMHasSnap	Boolean	TRUE if a VMware snapshot of the virtual machine is currently active.
VMHasVDSName	vmHasVdsName	Boolean	TRUE if the virtual machine has a display name that is valid for use as the host name.
hostName	hostName	Alphanumeric string	The virtual machine name that is derived from a reverse lookup of its IP address. Values are not case-sensitive.
vmIsConn	vmIsConn	Boolean	TRUE if the virtual machine is connected and available. For example: If a virtual machine's ESX server is down, that virtual machine is not connected.
vmVersion	vmVersion	Alphanumeric string	The VMware version of the virtual machine. Values are case-sensitive. For example: vmx-04, vmx-07, vmx-08.
VMXDatastore	vmxDatastore	Alphanumeric string	The name of the vmx datastore (sometimes called the vmx directory or configuration datastore). Values are case-sensitive. More information on the vmx datastore is available. See "NetBackup for VMware terminology" on page 21.
VMXDatastoreFolder	vmxDatastoreFolder	Alphanumeric string	The name of the folder that contains the vmx datastores. Values are not case-sensitive.

Table 8-15 Keywords in the **Field** dropdown *(continued)*

Field keyword	OData field keyword	Data type	Description
VMXDatastoreNFSHost	vmxDatastoreNfsHost	Alphanumeric string	The name of the vmx datastore's NFS host. Values are not case-sensitive.
VMXDatastoreNFSPath	vmxDatastoreNfsPath	Alphanumeric string	The folder structure that defines the path to a vmx NFS datastore. Use this option if the NFS host name of the datastore that you want to filter on is not unique in your environment. Values are not case-sensitive.
VMXDatastoreType	vmxDatastoreType	Alphanumeric string	The type of the vmx datastore. Values are NFS or VMFS. Values are not case-sensitive.
[vSphere custom attributes]	[vSphere custom attributes]	Alphanumeric string	The value of a custom attribute that is set in vSphere Client for one or more virtual machines. Note: In vSphere Client, the attribute must have a value for at least one virtual machine. The attribute type must be Virtual Machine. The values are case-sensitive. A disk exclusion based on custom attributes requires that you enter in NetBackup the credentials for the vCenter server or servers that host the VMs. ESXi server credentials are not sufficient. See “Add VMware servers” on page 66.

Operators

Table 8-16 describes the operators available in the **Operator** list.

Note: Use OData operators only when you build queries with the Advanced mode of the Query builder or with NetBackup APIs.

Table 8-16 Operators in the **Operator** list

Operator	OData operator	Description
NotEqual	ne	Matches any value that is not equal to the value in the Value field.
StartsWith	startswith	Matches the value in the Value field when it occurs at the start of a string. For example: If the Value entry is "box", StartsWith matches the string "box_car" but not "flatbox".
In	in	Matches any of the specified values in the Values field. For example: If the ESX servers in the Values field are "ESX01","ESX02,"ESX03", AnyOf matches any ESX server that has one of those names. If the names of your servers are not identical to any of the specified values, no match occurs. A server named ESX01A is not a match.
EndsWith	endswith	Matches the value in the Value field when it occurs at the end of a string. For example: If the Value entry is "dev", EndsWith matches the string "01dev" but not "01dev99", "devOP", or "Development_machine".
GreaterEqual	ge	Matches any value that is greater than or equal to the specified Value, according to the UTF-8 collating sequence.
Less	lt	Matches any value that is less than the specified Value, according to the UTF-8 collating sequence.
Equal	eq	Matches only the value that is specified in the Value field. For example: If the display name to search for is "VMtest27", Equal matches virtual machine names such as VMTest27 or vmtest27 or vmTEST27, and so forth. The name VMtest28 is not matched.
Greater	gt	Matches any value that is greater than the specified Value, according to the UTF-8 collating sequence.
Contains	contains	Matches the value in the Values field wherever that value occurs in the string. For example: If the Value entry is "dev", Contains matches strings such as "01dev", "01dev99", "devOP", and "Development_machine".
LessEqual	le	Matches any value that is less than or equal to the specified Value, according to the UTF-8 collating sequence.

Value

Table 8-17 describes the characters that can be entered in the **Value** field. The **Field** keyword determines case sensitivity.

Note: The character string you enter in the **Value** field must be enclosed in single quotes or double quotes.

Table 8-17 Characters you can enter for Value

Character types	String characters allowed
Alphanumerics	A to Z, a to z, 0 to 9, - (minus sign), and special characters. **Note:** Decimal numbers only.
Wildcards	* (asterisk) matches everything. For example: "*prod*" matches the string "prod" preceded or followed by any characters. ? (question mark) matches any single character. For example: "prod??" matches the string "prod" followed by any two characters.
Escape character	\ (backslash) escapes the wildcard or meta-character that follows it. For example: To search for a string that contains an asterisk (such as test*), enter "test*"
Quotation marks	Note: The characters you enter in Value must be enclosed in single or double quotes. To search for a string that contains quotation marks, either escape each quote (\") or enclose the entire string in the opposite type of quotes. For example: To search for a string that includes double quotes (such as "name"), enter '"name"' (enclosing it in single quotes) or "\"name\""

Test Query screen for VMware

This screen lists the virtual machines that NetBackup discovered in your virtual environment when you clicked **Test Query**. Later changes in the virtual environment may affect which virtual machines match the query rules. For example: if virtual machines are added, the test results may not be identical to the virtual machines that are selected for backup when the backup runs.

When the next backup runs from this policy, the following occur: NetBackup re-discovers virtual machines, consults the query rules, and backs up the virtual machines that match the rules.

The list of backed up virtual machines is saved but the virtual machines are not displayed in the policy's **Clients** tab. You can use the Activity Monitor to view the virtual machine jobs.

Note: An alternative to the Test Query screen is the `nbdiscover` command. For more information, see the *NetBackup Commands Reference Guide*.

See "Using the Activity monitor to monitor virtual machine backups" on page 222.

The **Test Query** function runs in the background. You can continue to configure the policy while the test runs. Any changes you make in the Query Builder however are not included in the currently running test. You must re-initiate the test to see the results of your Query Builder changes.

Table 8-18

Field	Description
Test query for policy	Lists the rules in the Query Builder that were used in this test. The rules are specified in the Query Builder on the policy **Clients** tab.
Test Query Results	**VM Name**: Shows the display name of all discovered virtual machines. **Selection**: Lists the virtual machines that were discovered, as follows: ■ **INCLUDED**: The virtual machine matches the rules in the query. ■ **EXCLUDED**: The virtual machine does not match the rules in the query. ■ **FAILED**: The virtual machine cannot be selected for backup because of a host name problem or other error. Also, the query cannot exclude the virtual machine. An explanation appears at the bottom of the **Test Query** screen. For example: `VM does not have a host name to use as a client name, display name =` See "Test Query: Failed virtual machines" on page 154. The operator IsSet can be used to filter out such virtual machines. More information is available on IsSet. See Table 8-16 on page 151. See "The IsSet operator in queries" on page 132.
Included: **Excluded:** **Failed:**	Gives a tally of how many virtual machines were included, excluded, or failed in the test.

See "Using the Activity monitor to monitor virtual machine backups" on page 222.

See "About automatic virtual machine selection for NetBackup for VMware" on page 115.

See "Configure automatic virtual machine selection" on page 125.

Test Query: Failed virtual machines

If the query rules cannot exclude a virtual machine, and that virtual machine cannot be selected for backup, it is marked as FAILED. The virtual machine is listed as not run in the job details log.

For example: the virtual machine does not have the type of name specified by the **Primary VM identifier** parameter (such as host name or display name). Or the virtual machine name contains invalid characters. In any case, a virtual machine that is listed as FAILED should be investigated: it may be one that you want to back up.

To see the reason for the failure, click on the virtual machine in the Test query results. An explanation appears at the bottom of the screen.

You can fix this problem in a couple of ways:

- Use vSphere Client to configure a host name for the virtual machine.
- To exclude the virtual machines that have no host name, construct a query with the IsSet operator.

See "The IsSet operator in queries" on page 132.

Effect of Primary VM identifier parameter on Selection column in Test Query results

The NetBackup policy's **Primary VM identifier** parameter tells NetBackup how to identify virtual machines. For example, if the parameter is set to **VM hostname**, NetBackup identifies virtual machines by their host names. If they do not have a host name, the policy cannot back them up.

The **Primary VM identifier** parameter has a direct effect on the query test results. Note that for each virtual machine, the query test result is one of three possibilities: INCLUDED, EXCLUDED, or FAILED.

If NetBackup cannot identify a virtual machine according to the **Primary VM identifier** parameter, one of two test results can occur:

- If the virtual machine is filtered out by the query rules, it is listed as EXCLUDED.
- If the virtual machine is not filtered out by the query rules, it is listed as FAILED. The following table gives the test query results from example combinations of the **Primary VM identifier** parameter and a query rule.

Table 8-19 Effect of Primary VM identifier parameter and query rules on test query results

Primary VM identifier setting on VMware policy tab	Query rule in Query Builder	OData query rule in Query Builder *	Test query result
VM hostname	VMHostName Contains "VM"	contains(displayName, 'VM')	INCLUDED: Any virtual machines with a host name that contains "VM". Since the Primary VM identifier parameter tells NetBackup to select the virtual machine by host name, it can back up the virtual machines. EXCLUDED: All other virtual machines.
VM hostname	Displayname Contains "VM"	contains(displayName, 'VM') and hostName ne null	INCLUDED: Any virtual machines that have a host name and that have a display name that contains "VM". EXCLUDED: Any virtual machines that have a host name, but that do not have a display name containing "VM". FAILED: Any virtual machines that do not have a host name. Since the **Primary VM identifier** parameter is set to VM hostname, NetBackup cannot select the virtual machine for backup.
VM hostname	Displayname Contains "VM" AND VMHostName IsSet	contains(displayName, 'VM') and hostName ne null	INCLUDED: Any virtual machines that have a host name and that have a display name that contains "VM". EXCLUDED: All other virtual machines. The IsSet rule means that if a virtual machine does not have a host name, it is excluded.

Table 8-19 Effect of Primary VM identifier parameter and query rules on test query results *(continued)*

Primary VM identifier setting on VMware policy tab	Query rule in Query Builder	OData query rule in Query Builder *	Test query result
VM hostname	Displayname Contains "VM" AND VMHostName IsSet OR Annotation Contains "test" AND NOT VMHostName IsSet	contains(displayName, 'VM') and hostName ne null or contains(annotation, 'test') and not (hostName ne null)	INCLUDED: ■ Any virtual machines that have a host name and that have a display name that contains "VM". ■ Any virtual machines without a host name that have an annotation that contains "test". EXCLUDED: All other virtual machines.
VM display name	Displayname Contains "VM"	contains(displayName, 'VM')	INCLUDED: Any virtual machines with the display names that contain "VM". Since the Primary VM identifier parameter tells NetBackup to select the virtual machine by display name, it can back up the virtual machines. EXCLUDED: All other virtual machines.
VM display name	VMHostName Contains "VM"	contains(hostName, 'VM')	INCLUDED: Any virtual machines that have a display name and that have a host name that contains "VM". EXCLUDED: Any virtual machines that have a display name, but that do not have a host name containing "VM". FAILED: Any virtual machines that do not have a display name. Since the **Primary VM identifier** parameter is set to VM display name, NetBackup cannot select those virtual machines for backup.

Table 8-19 Effect of Primary VM identifier parameter and query rules on test query results *(continued)*

Primary VM identifier setting on VMware policy tab	Query rule in Query Builder	OData query rule in Query Builder *	Test query result

* Use OData keywords only when you build queries with the NetBackup web UI.

Effect of Primary VM identifier parameter on VM Name column in Test query results

The policy's **Primary VM identifier** parameter affects the type of virtual machine name that appears in the **VM Name** column of the Test Query screen, as follows:

- If a virtual machine is EXCLUDED or FAILED, it is listed according to its virtual machine display name. The **Primary VM identifier** parameter does not matter.
- But if a virtual machine is listed as INCLUDED, note: The name that appears under **VM Name** is the type of name that is specified on the **Primary VM identifier** parameter.
  For example: If the **Primary VM identifier** parameter is VM hostname, the included virtual machine is listed according to its host name. Even if the query rule specified Display name (such as `Displayname Equal "vm1"`), the virtual machine appears on the Test Query screen by its host name.

See "Primary VM identifier options (VMware)" on page 93.

Refreshing the display of virtual environment changes in the Query Builder

By default, NetBackup waits one hour before the policy Query Builder detects changes in the virtual environment. Until one hour has passed, the Query Builder does not detect the changes when you click the "Load values" folder icon next to the **Value** field. To make the changes immediately available to the **Value** field, use the following procedure to refresh the display.

Note: The Query Builder's **Reuse VM selection query results for** option does not affect the display of virtual environment changes in the Query Builder. The reuse option determines how long NetBackup reuses the current backup list for future executions of the policy.

To refresh the Query Builder's view of the virtual environment (Windows):

1. On the Windows desktop of the local host, click **Start > Run** and enter `regedit`.
2. Make a backup of the current registry (**File > Export**).
3. Go to **HKEY_LOCAL_MACHINE > SOFTWARE > Veritas > NetBackup > CurrentVersion > Config** and create a key that is called `BACKUP`.
4. Under `BACKUP`, create a new DWORD that is called `xmlCacheLimit`.
5. Set this DWORD to the number of seconds for the refresh.

   A value of 15 allows the Query Builder to be refreshed after 15 seconds.
6. If the policy editor is open, close it and reopen it.

To refresh the Query Builder's view of the virtual environment (Linux):

1. On the Linux desktop of the local host, create (or open) the following file:

   `/usr/openv/netbackup/virtualization.conf`
2. Enter a new `dword` line under `[BACKUP]` to set the number of seconds for the refresh. For example:

   ```
   [BACKUP]
   "xmlCacheLimit"=dword:15
   ```

   This example allows the Query Builder to be refreshed after 15 seconds.

   Note: If the file already contains a `[BACKUP]` line, do not add another `[BACKUP]` line. Any other lines that already exist under `[BACKUP]` should remain as they are.

   ---
3. Save the file.
4. If the policy editor is open, close it and reopen it.

Reducing the time required for VM discovery in a large VMware environment

NetBackup VMware intelligent policies use query rules to automatically search and filter the vSphere environment. By default, the query rules search all the VMware servers in your environment. If the environment contains many VMware servers with many virtual machines, VM discovery may take a long time. You can speed up VM discovery by limiting the search to specific VMware servers or virtual machines.

The following is an example of a policy Query Builder rule that searches all VMware servers and all virtual machines:

```
vmware:/?filter=Displayname Contains "vm1"
```

To limit the search to particular servers or virtual machines, insert an additional expression in the Query Builder rule as explained in the following procedure.

To reduce the time required for VM discovery in a VMware environment

1. In the NetBackup web UI, open the VMware Intelligent Policy.
2. On the **Clients** tab of the policy, make sure **Select automatically through VMware intelligent policy query** is selected.
3. Under the **Query builder**, click **Advanced mode**.
4. Create one or more rules to search for VMs in specific VMware servers or for specific virtual machines.

Be aware that to create two or more rules, you must be in **Advanced mode** not **Basic mode**. Additionally, each query rule must begin on its own line.

You can use the types of rules shown:

- To search for VMs in a particular VMware server

  vmware://*VMware_server*?filter=*filter*
- To search for a specific virtual machine on a particular VMware server

  vmware://*VMware_server*/vm/*virtual_machine_instance_uuid*

Chapter 9

Use Accelerator to back up virtual machines

This chapter includes the following topics:

About the NetBackup Accelerator for virtual machines

NetBackup Accelerator reduces the backup time for VMware backups. NetBackup uses VMware Changed Block Tracking (CBT) to identify the changes that were made within a virtual machine. Only the changed data blocks are sent to the NetBackup media server, to significantly reduce the I/O and backup time. The media

server combines the new data with previous backup data and produces a traditional full NetBackup image that includes the complete virtual machine files.

Note: Accelerator is most appropriate for virtual machine data that does not experience a high rate of change.

Accelerator has the following benefits:

- Performs the full backups faster than traditional backup. Creates a compact backup stream that uses less network bandwidth between the backup host and the server.
  Accelerator sends only changed data blocks for the backup. NetBackup then creates a full traditional NetBackup image that includes the changed block data.
- Accelerator backups support Granular Recovery Technology (GRT) for restoring Exchange, SQL, and SharePoint applications (using a full schedule only).
- Accelerator backups (full and incremental) support instant recovery of virtual machines.
- If the **Enable file recovery from VM backup** option on the policy **VMware** tab is enabled, you can restore individual files from the backup (full or incremental).
- Reduces the I/O on the backup host.
- Reduces the CPU load on the backup host.

Accelerator: full vs. incremental schedules

NetBackup Accelerator supports full and incremental backups.

Note: After an initial full backup, Accelerator backups with a full schedule have about the same effect on I/O and performance as traditional incremental backups. The NetBackup catalog however includes all catalog references that would be made if the backup was a traditional (non-Accelerator) full.

Note: After upgrade from any of the previous release, If the customer has configured accelerator enabled Hyper-v policy, then optimization will be lost only for first backup.

For virtual machine restore, note the following about full vs. incremental backups with Accelerator:

- For applications (Exchange, SQL, and SharePoint), NetBackup Accelerator supports Granular Recovery Technology (GRT) restores from full backups only.

- For any other kind of virtual machine restore, the Accelerator supports full backups and incremental backups.

How the NetBackup Accelerator works with virtual machines

To enable acceleration of virtual machine backups, click **Use Accelerator** on the policy **Attributes** tab.

See "Configure a VMware policy" on page 87.

The NetBackup Accelerator creates the backup stream and backup image for each virtual machine as follows:

- If the virtual machine has no previous backup, NetBackup performs a full backup and uses VMware Changed Block Tracking to track the data in use for each vmdk.
- At the next backup, NetBackup identifies data that has changed since the previous backup. Only changed blocks and the header information are included in the backup, to create a full virtual disk backup.
- The backup host sends to the media server a tar backup stream that consists of the following: The virtual machine's changed blocks, and the previous backup ID and data extents (block offset and size) of the unchanged blocks.
- The media server reads the virtual machine's changed blocks, the backup ID, and information about the data extents of the unchanged blocks. From the backup ID and data extents, the media server locates the rest of the virtual machine's data in existing backups.
- The media server directs the storage server to create a new full image that consists of the following: The newly changed blocks, and the existing unchanged blocks that reside on the storage server. The storage server may not write the existing blocks but rather link them to the image.

Accelerator notes and requirements for virtual machines

Note the following about Accelerator for virtual machines:

- Accelerator for virtual machines uses VMware Changed Block Tracking (CBT) to identify the changes that were made within a virtual machine.
  VMware CBT may occasionally reset tracking of file changes, such as after a power failure or hard shutdown. In that case, for the next backup NetBackup

reads all the data from the vmdk files and the backup takes longer than expected. If deduplication is enabled, the deduplication rate is lower than expected.
For more information on CBT, see the following VMware article:

- Supports the disk storage units that have the following storage destinations:
  - Cloud storage. Storage that a supported cloud storage vendor provides.
  - NetBackup Media Server Deduplication Pool. In addition to NetBackup media servers, NetBackup 5200 series appliances support Media Server Deduplication Pool storage.
  - Qualified third-party OpenStorage devices.

  To verify that your storage unit supports Accelerator, refer to the NetBackup hardware compatibility list for the currently supported OST vendors:
- It is recommended that you do not enable **Expire after copy** retention for any storage units that are used with storage lifecycle policies (SLP) in combination with Accelerator. The **Expire after copy** retention can cause images to expire while the backup runs. To synthesize a new full backup, the SLP backup needs the previous backup image. If the previous image expires during the backup, the backup fails.
- Update the NetBackup device mapping files if needed.
  The NetBackup device mapping files contain all storage device types that NetBackup can use. To add support for the new devices or upgraded devices that support Accelerator, download the current device mapping files from the Veritas Technical Support website.
  See the *NetBackup Administrator's Guide Volume I* for information on the device mapping files and how to download them.
- Storage unit groups are supported only if the storage unit selection in the group is Failover.
- Supports the full backups and incremental backups. Every Accelerator backup (from a full schedule or incremental schedule) results in a complete image of the virtual machine.
- You can use incremental backups (cumulative or differential) as follows: To reduce the file-mapping overhead and to reduce the number of files that are recorded in the NetBackup catalog. Cumulative backups may involve more file-mapping because they do not use the random indexing method to determine which files have changed. In some cases, differential backups may be faster than cumulative backups.
- If a backup of the virtual machine does not exist, NetBackup performs a full backup. On the backup host it also accesses the VMware CBT information. This initial backup occurs at the speed of a normal (non-accelerated) full backup.

Subsequent Accelerator backups of the virtual machine use VMware Changed Block Tracking to accelerate the backup.

Note: When you first enable a VMware policy to use Accelerator, the next backup (whether full or incremental) is in effect a full backup: It backs up all the virtual machine files that are selected in the policy. If that backup is an incremental, it may not complete within the backup window.

- If the storage unit that is associated with the policy cannot be validated when you create the policy, note: The storage unit is validated later when the backup job begins. If Accelerator does not support the storage unit, the backup fails. In the `bpbrm` log, a message appears that is similar to one of the following:

```
Storage server %s, type %s, doesn't support image include.

Storage server type %s, doesn't support accelerator backup.
```

- Accelerator requires the storage to have the `OptimizedImage` attribute enabled. See "Accelerator requires the `OptimizedImage` attribute" on page 165.
- Because of a VMware restriction, BLIB is not supported for VMware templates. As a result, NetBackup Accelerator cannot be used to back up VMware virtual machine templates.

Accelerator forced rescan for virtual machines (schedule attribute)

The accelerator for virtual machines uses Changed Block Tracking (CBT) technology from VMware to identify changed blocks. NetBackup requires the changed blocks when it creates a full virtual machine (synthesized) image. NetBackup is therefore dependent on VMware CBT for correctly identifying changed blocks. To protect against any potential omissions by underlying VMware CBT using timestamps, the **Accelerator forced rescan** option conducts the backup by collecting all in-use blocks as reported by VMware CBT.

When **Accelerator forced rescan** is used, all the data on the virtual machine is backed up. This backup is similar to the first Accelerator backup for a policy. For the forced rescan job, the optimization percentage for Accelerator is 0. The duration of the backup is similar to a non-Accelerator full backup.

Note: Under normal operations, an Accelerator forced rescan schedule is not necessary for VMware backups. It can be used to enforce a new baseline backup in case VMware CBT issues are discovered. Engaging with Cohesity Support is recommended in such situations.

Accelerator forced rescan is unavailable if the **Use Accelerator** option on the **Attributes** tab is not selected.

Accelerator requires the `OptimizedImage` attribute

Accelerator requires that the storage has the `OptimizedImage` attribute enabled.

To ensure that your storage is configured properly, see the documentation for your storage option:

- NetBackup Media Server Deduplication Pool.
  The `OptimizedImage` attribute is enabled by default. If you created the storage servers and pools in an earlier release, you must configure them for `OptimizedImage`.
  See the *NetBackup Deduplication Guide*.
- Backups to a third-party disk appliance.
  The storage device must support the `OptimizedImage` attribute.
  See the *NetBackup OpenStorage Solutions Guide for Disk*.
- Cloud storage that NetBackup supports.
  See the *NetBackup Cloud Administrator's Guide*.

Accelerator backups and the NetBackup catalog

Use of Accelerator does not affect the size of the NetBackup catalog. A full backup with Accelerator generates the same catalog size as a full backup of the same data without Accelerator. The same is true of incremental backups: use of Accelerator does not require more catalog space than the same backup without Accelerator.

However, if you enable **Calculate file hash** option on the web UI, the NetBackup catalog is expected to increase by 20% or more. This option is used to save the file hash information to the NetBackup catalog. See the *NetBackup Web UI Administrator's Guide* for more information.

A potential catalog effect does exist, depending on how often you use Accelerator with full backups. A full backup with Accelerator completes faster than a normal full. It may therefore be tempting to replace your incremental backups with Accelerator full backups. Note: Since a full backup requires more catalog space than an incremental, replacing incrementals with fulls increases the catalog size.

When changing your incrementals to fulls, you must weigh the advantage of Accelerator fulls against the greater catalog space that fulls require compared to incrementals.

Accelerator messages in the backup job details log

When a virtual machine is first backed up, Accelerator is not used for that backup. The following messages appear in the job details log:

```
7/25/2012 4:45:35 PM - Info bpbrm(pid=6192) There is no complete
 backup image match with track journal, a regular full backup will
 be performed

...

7/25/2012 4:53:22 PM - Info bpbkar32(pid=5624) accelerator sent
5844728320 bytes out of 5844726784 bytes to server, optimization 0.0%
```

When subsequent backups of the virtual machine use Accelerator, the following messages appear in the job details log:

```
7/27/2012 4:40:01 AM - Info bpbrm(pid=412) accelerator enabled

...

7/27/2012 4:43:07 AM - Info bpbkar32(pid=4636) accelerator sent
74764288 bytes out of 5953504256 bytes to server, optimization 98.7%
```

This message is a key trace for Accelerator. In this example Accelerator was successful at reducing the backup data by 98.7%.

About reporting the amount of Accelerator backup data that was transferred over the network

For Accelerator backup reporting, several NetBackup commands can report the amount of data that is transferred over the network for each Accelerator backup. The amount of transferred data is often much less than the size of the Accelerator backup image.

For each Accelerator backup, NetBackup combines the client's (or VM's) changed blocks with the unchanged data from previous backups to synthesize a backup image. However, NetBackup sends only the changed data over the network when the backup occurs. The resulting backup image may be much larger than the amount

of backup data that travels the network. For backup reporting, it may be important to distinguish between the backup image size and the amount of data that was transferred over the network.

For Accelerator backups, the network-transferred data can appear in the output of the following NetBackup commands: `bpdbjobs`, `bpimagelist`, and `bpclimagelist`.

Table 9-1 lists the default location of these commands.

Table 9-1 Default location of `bpdbjobs`, `bpimagelist`, and `bpclimagelist`

Command	Default location
`bpdbjobs`, `bpimagelist`	Windows: *`install_path`*`\NetBackup\bin\admincmd\` UNIX, Linux `/usr/openv/netbackup/bin/admincmd/`
`bpclimagelist`	Windows: *`install_path`*`\NetBackup\bin\` UNIX, Linux `/usr/openv/netbackup/bin/`

The following example uses the `bpimagelist` command to show the results of a backup of `acmevm2`:

```
bpimagelist -backupid acmevm2
```

Example output:

```
IMAGE acmevm2 0 0 12 acmevm2 accl_vmware 40 *NULL* root f 0 9 14344
79628 558 2147483647 0 0 7799632 28196 1 2 0 accl_vmware_1434479628_FULL.f *NULL
* *NULL* 0 1 0 0 0 *NULL* 0 0 1 0 0 1434479628 1434479628 *NULL* 0 0 0 *NULL* 9
0 0 3398732 0 0 *NULL* *NULL* 0 1434479620 0 0 *NULL* *NULL* 0 0 0 225792
HISTO 0 0 0 0 0 0 0 0 0 0
FRAG 1 -1 3319 76 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 102
4 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 2147483647
0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
FRAG 1 1 7796313 0 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 10
28 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 214748364
7 0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
```

In this example, the backup image size in kilobytes is 7799632, and the amount of data that was transferred over the network is 225792.

You can use the following commands to show the amount of data that was transferred over the network for an Accelerator backup.

bpimagelist

```
bpimagelist -backupid backup_id [-l | -L | -json | -json_compact]
```

Brackets [ ] indicate optional elements, and the vertical bars | indicate that you can choose only one of the options within the brackets.

Table 9-2 describes how the network-transferred data field appears in the `bpimagelist` output.

Table 9-2 The **bpimagelist** options that show the amount of network-transferred data for Accelerator backups

bpimagelist option	How the network-transferred data field appears
No option	The field is unlabeled. For example: `225792` See the `bpimagelist` example output earlier in this topic.
`-l`	The field is unlabeled (same as no option). For example: `225792`
`-L`	The field is labeled. For example: `Kilobytes Data Transferred: 225792`
`-json`	The field is labeled. For example: `"kilobytes_data_transferred": 225792,`
`-json_compact`	The field is labeled. For example: `"kilobytes_data_transferred":225792,`

bpdbjobs

```
bpdbjobs -jobid job_id -report -most_columns
```

or

```
bpdbjobs -jobid job_id -report -all_columns
```

The network-transferred data field appears at the end of the output.

bpclimagelist

```
bpclimagelist -client client_name
```

This command can only show the network-transferred data in the field that normally shows the Accelerator backup image size. To show the network-transferred data with this command, you must configure a NetBackup setting:

See “Replacing the Accelerator image size with the network-transferred data in NetBackup command output” on page 169.

Additional details on these commands are available in the *NetBackup Commands Reference Guide*.

Replacing the Accelerator image size with the network-transferred data in NetBackup command output

You can configure the output of `bpimagelist`, `bpdbjobs`, and `bpclimagelist` to show the amount of Accelerator backup data that was transferred over the network instead of the backup image size.

The following is the default `bpimagelist` output that shows the Accelerator image size (see the circled value 7799632). The amount of network-transferred data appears farther down in the output (225792):

```
IMAGE acmevm2 0 0 12 acmevm2 accl_vmware 40 *NULL* root f 0 9 14344
79628 558 2147483647 0 0 7799632 28196 1 2 0 accl_vmware_1434479628_FULL.f *NULL
* *NULL* 0 1 0 0 0 *NULL* 0 0 1 0 0 1434479628 1434479628 *NULL* 0 0 0 *NULL* 9
0 0 3398732 0 0 *NULL* *NULL* 0 1434479620 0 0 *NULL* *NULL* 0 0 0 225792
HISTO 0 0 0 0 0 0 0 0 0 0
FRAG 1 -1 3319 76 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 102
4 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 2147483647
0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
FRAG 1 1 7796313 0 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 10
28 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 214748364
7 0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
```

You can configure NetBackup command output to show the network-transferred data in the image size field. In the output, the image size value is replaced with the network-transferred data value (see the following example). A script that reads the image size from the command output now reads the amount of network-transferred data.

In the following `bpimagelist` output, the image size field shows the network-transferred data (225792):

```
IMAGE acmevm2 0 0 12 acmevm2 accl_vmware 40 *NULL* root f 0 9 14344
79628 558 2147483647 0 0 225792 28196 1 2 0 accl_vmware_1434479628_FULL.f *NULL
* *NULL* 0 1 0 0 0 *NULL* 0 0 1 0 0 1434479628 1434479628 *NULL* 0 0 0 *NULL* 9
0 0 3398732 0 0 *NULL* *NULL* 0 1434479620 0 0 *NULL* *NULL* 0 0 0 225792
HISTO 0 0 0 0 0 0 0 0 0 0
FRAG 1 -1 3319 76 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 102
4 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 2147483647
0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
FRAG 1 1 7796313 0 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 10
28 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 214748364
7 0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
```

Note: The same change occurs in the labeled output of the commands (such as with the `-L` option of `bpimagelist`). For example, the `Kilobytes` field shows the transferred data value (225792 in the example) rather than the Accelerator backup image size.

To enable the reporting of network-transferred data in the Accelerator image size field of bpimagelist, bpdbjobs, and bpclimagelist

- Use the `bpsetconfig` command to enable the output change.

  To enable this change for the `bpclimagelist` command, enter the `bpsetconfig` command on the primary server. To enable this change for `bpimagelist` or `bpdbjobs`, enter the `bpsetconfig` command on the server where you intend to run `bpimagelist` or `bpdbjobs`.

  Refer to Table 9-3 for the `bpsetconfig` command to use based on the type of Accelerator backup that you want to report on.

Table 9-3 To enable the reporting of network-transferred data in the Accelerator image size field of **bpimagelist**, **bpclimagelist**, or **bpdbjobs** output

Type of backup to report on	Enter this command
Incremental VMware Accelerator backups	Windows `echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_ FOR_ACCL_INC_VMWARE	install_path\NetBackup\bin\admincmd\bpsetconfig` UNIX, Linux `echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_ FOR_ACCL_INC_VMWARE"	/usr/openv/netbackup/bin/admincmd/bpsetconfig`
All VMware Accelerator backups (full and incremental)	Windows `echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_ FOR_ACCL_ALL_VMWARE	install_path\NetBackup\bin\admincmd\bpsetconfig` UNIX, Linux `echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_ FOR_ACCL_ALL_VMWARE"	/usr/openv/netbackup/bin/admincmd/bpsetconfig`
Incremental Accelerator virtual machine backups (VMware and Hyper-V)	Windows `echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_ FOR_ACCL_INC_VIRTUAL	install_path\NetBackup\bin\admincmd\bpsetconfig` UNIX, Linux `echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_ FOR_ACCL_INC_VIRTUAL"	/usr/openv/netbackup/bin/admincmd/bpsetconfig`
All Accelerator virtual machine backups (VMware and Hyper-V, full and incremental)	Windows `echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_ FOR_ACCL_ALL_VIRTUAL	install_path\NetBackup\bin\admincmd\bpsetconfig` UNIX, Linux `echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_ FOR_ACCL_ALL_VIRTUAL"	/usr/openv/netbackup/bin/admincmd/bpsetconfig`

Table 9-3 To enable the reporting of network-transferred data in the Accelerator image size field of **bpimagelist**, **bpclimagelist**, or **bpdbjobs** output *(continued)*

Type of backup to report on	Enter this command

All incremental Accelerator backups (physical clients and virtual machines)	Windows
	`echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_INC_ALL | install_path\NetBackup\bin\admincmd\bpsetconfig`
	UNIX, Linux
	`echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_INC_ALL" | /usr/openv/netbackup/bin/admincmd/bpsetconfig`

All Accelerator backups (full and incremental, physical clients and virtual machines)	Windows
	`echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_ALL_ALL | install_path\NetBackup\bin\admincmd\bpsetconfig`
	UNIX, Linux
	`echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_ALL_ALL" | /usr/openv/netbackup/bin/admincmd/bpsetconfig`

To reset the command output to the default setting

- To disable the reporting of network-transferred data in the Accelerator image size field (return to default), enter the following:

  Windows

  ```
  echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED =
  REPLACE_IMAGE_SIZE_DISABLED |
  install_path\NetBackup\bin\admincmd\bpsetconfig
  ```

  UNIX, Linux

  ```
  echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED =
  REPLACE_IMAGE_SIZE_DISABLED" |
  /usr/openv/netbackup/bin/admincmd/bpsetconfig
  ```

Chapter 10

Configuring protection plans for VMware

This chapter includes the following topics:

Protect VMs or intelligent VM groups

Use the following procedure to subscribe an asset (VMs or intelligent VM groups) to a protection plan. When you subscribe an asset to a protection plan, you assign predefined backup settings to the asset.

Note: Protection plans are not supported for VMware Cloud Director VMs.

To protect VMs or VM groups

1. On the left, click **Workloads > VMware**.
2. On the **Virtual machines** tab or **Intelligent VM groups** tab, click the box for the VM or the VM group and click **Add protection**.
3. Select a protection plan and click **Next**.
4. Adjust any settings as necessary.
   - Change the backup start window.
     See “Schedules” on page 174.

- **Backup options** and **Advanced** options.
  See "Backup options and Advanced options" on page 174.

5 Click **Protect**.

The results of your choices appear under **Virtual machines** or **Intelligent VM groups**.

Schedules

The following schedule settings are included in a protection plan.

Note that when you customize a protection plan for an asset, you can only edit the following schedule settings:

- Start window

Table 10-1 Schedule options for protection plans

Option	Description
Backup type	The type of backup that the schedule controls.
Recurrence (frequency)	How frequently or when to run the backup.
Keep for (retention)	How long to keep the files that were backed up by the schedule.
Replicate this backup	Replicates the snapshot to another volume.
Duplicate a copy immediately to long-term retention	Immediately after the schedule is created, a copy is duplicated to the media that is selected for long-term storage.
Start window	On this tab, set the window during which a backup can start.

Backup options and Advanced options

The user can adjust the following settings when subscribing to a protection plan.

Backup options

Table 10-2 Backup options for protection plans

Option	Description
Select server or host to use for backups	The host that performs backups on behalf of the virtual machines. Users can choose **Automatic** to have NetBackup pick the media server, based on the storage unit. Or, the user can select another host from the list. These hosts are other media servers in the environment or hosts that are configured as an access host.
If a snapshot exists, perform the following action	Specifies the action that NetBackup takes when a snapshot is discovered before NetBackup creates a new snapshot for the virtual machine backup. For example, users can choose to stop a backup if any snapshots exist. If snapshots are not automatically deleted, the performance of the virtual machine may eventually decline. Undeleted snapshots can cause restore failures due to lack of disk space.
Exclude selected virtual disks from backups	Specifies the virtual disks to exclude from backups. See "Exclude disks from backups" on page 176.

Advanced options

Table 10-3 Advanced options for protection plans

Option	Description
Enable virtual machine quiesce	By default, I/O on the virtual machine is quiesced before NetBackup creates the snapshot. In the majority of cases, you should use this default. Without quiescing file activity, data consistency in the snapshot cannot be guaranteed. If you disable the quiesce, you must analyze the backup data for consistency.
Allow the restore of application data from virtual machine backups	This option allows users to restore application data from full backups of the virtual machine. If this option is disabled, you can recover the entire virtual machine from the backup, but you cannot recover the databases individually. Note that in NetBackup 8.3 or earlier, application data for Microsoft Exchange Server or Microsoft SharePoint Server must be restored with the NetBackup Backup, Archive, and Restore interface. Data for Microsoft SQL Server must be restored with the NetBackup MS SQL Client. See the documentation for your NetBackup database agent for more details.
Transport mode	Specifies the transport mode to use for backups or how to read the data from the datastore. For more information on transport modes, see the vendor documentation for your virtualization environment.
Snapshot retry options	See "Snapshot retry options" on page 176.

Exclude disks from backups

Excluding virtual disks can reduce the size of the backup, but use these options carefully. They are intended only for the virtual machines that have multiple virtual disks.

Table 10-4 Options for excluding virtual disks

Exclude option	Description
All boot disks	Consider this option if you have another means of recreating the boot disk. The virtual machine's boot disk is not included in the backup. Any other disks are backed up. **Note:** Data files are available in the restored data disks. However, you cannot start a virtual machine that is restored from this backup.
All data disks	Consider this option only if you have a separate protection plan that backs up the data disks. The virtual machine's data disks are not included in the backup. Only the boot disk is backed up. **Note:** When the virtual machine is restored from the backup, the virtual machine data for the data disk may be missing or incomplete.
Exclude disks based on a custom attribute	Use this option to allow the VMware administrator to use a custom attribute to control which disks are excluded from backups. The attribute must have comma-separated values of device controllers for the disks to be excluded. For example: `scsi0-0,ide0-0,sata0-0,nvme0-0`. The default value for this attribute is `NB_DISK_EXCLUDE_DISK`. Or, you can choose your own value. If you add disks to the custom attribute value between any differential backups, those disks are excluded from the next backup. The VMware administrator must use a VMware interface to apply the attribute to the disks to exclude. See the NetBackup Plug-in for VMware vSphere Web Client Guide or the NetBackup Plug-in for VMware vSphere Client (HTML5) Guide.
Specific disks to be excluded	Use this option to exclude a specific disk by the disk type, controller, and LUN that represent the virtual device node of the disk. Click **Add** to specify additional disks. If you add controllers between any differential backups, their disks are excluded from the next backup.

Snapshot retry options

For most environments, the default values for the snapshot retry options are appropriate. It may be helpful to adjust these settings based on the size of the virtual machine and the processing load on the VMware server.

Table 10-5 Snapshot retry options

Option	Description
Maximum number of times to retry a snapshot	The number of times the snapshot is retried.
Maximum length of time to complete a snapshot	The time, in minutes, to allow the snapshot operation to complete. If snapshots do not complete, set this option to a specific period to force a time-out. Use the **Maximum length of time to wait before a snapshot is retried** setting to retry the snapshot at a later time.
Maximum length of time to wait before a snapshot is retried	The time to wait (in seconds) before the snapshot is retried.

Customize protection settings for a VMware asset

You can customize certain settings for a protection plan, including the schedule backup window and other options.

- See “Schedules” on page 174.
- See “Backup options and Advanced options” on page 174.

To customize protection settings for a VMware asset

1. On the left, click **Workloads > VMware**.
2. Do one of the following:

Edit the settings for a VM	■ On the **Virtual machines** tab, click on the VM that you want to edit.
Edit the settings for an intelligent group	■ On the **Intelligent VM groups** tab, click on the group that you want to edit.

3. Click **Customize protection > Continue**.
4. Adjust any of the following settings:
   - The backup start window.
     See “Schedules” on page 174.
   - **Backup options** and **Advanced** options.
     See “Backup options and Advanced options” on page 174.
5. Click **Protect**.

Remove protection from VMs or intelligent VM groups

You can unsubscribe VMs or intelligent VM groups from a protection plan. When the asset is unsubscribed, backups are no longer performed.

Note: When you unsubscribe an asset from a protection plan, there is a possibility that the asset displays **Classic policy** in the web UI. This situation can happen when an asset is subscribed to a protection plan and a backup runs for that asset. Then the asset is unsubscribed from the protection plan while it has a valid backup image. The web UI displays **Classic policy**, but there may or may not be an active policy protecting the asset.

To remove protection from a VM or intelligent VM group

1. On the left, click **Workloads > VMware**.
2. On the **Virtual machines** tab or **Intelligent VM groups** tab, click the VM or the intelligent VM group.
3. Click **Remove protection > Yes**.

   Under **Virtual machines** or **Intelligent VM groups**, the asset is listed as Not protected.

View the protection status of VMs or intelligent VM groups

You can view the protections plans that are used to protect VMs or intelligent VM groups.

To view the protection status of VMs or intelligent VM groups

1. On the left, click **Workloads > VMware**.
2. Select the **Virtual machines** tab or **Intelligent VM groups** tab, as appropriate.

   Note: Sorting on assets across asset types, that is, without the Asset Type filter, returns results grouped by asset types (Virtual Machine and Intelligent VM groups) and sorted within each asset type.

   ---

3 Click the VM or the intelligent VM group.

The **Protection** tab shows the details of the plans that the asset is subscribed to.

Note: If the asset has been backed up, but Status indicates it has not, see the following information.

See "Troubleshooting the status for a newly discovered VM" on page 330.

---

4 If the asset is not protected, click **Add protection** to select a protection plan.

See "Protect VMs or intelligent VM groups" on page 173.

Chapter 11

Malware scan

This chapter includes the following topics:

Assets by workload type

This section describes the procedure for scanning VMware VM assets for malware.

This section describes the procedure for scanning VMware, Universal shares, Kubernetes, Nutanix and Cloud VM assets for malware.

To scan the supported assets for malware, perform the following:

1. On left, select the supported workload under **Workloads**.
2. Select the resource which has backups completed.

   For example, VMware, Universal shares, Kubernetes, Nutanix and Cloud VM

   For example, VMware

   For example, Nutanix AHV
3. Select **Actions > Scan for malware**.
4. On the **Malware scan** page, perform the following:
   - Select the date range for the scan by selecting **Start date/time** and **End date/time**.
   - Select **Scanner host pool**
   - From the **Current infection status** list select one of the following:
     - **Not scanned**
     - **Not infected**
     - **Infection detected by malware scan**

- **Infection detected by file hash search**
- **All**

5 Click **Scan for malware**.

Note: The malware scanner host can initiate a scan of three images at the same time.

---

6 After the scan starts, you can see the **Scan status** on **Malware detection**, the following fields are visible:

- **Not scanned**
- **Not infected**
- **Infected**
- **Failed**

Note: Any backup images that fail validation are ignored.

- **In progress**
- **Pending**

Chapter 12

Instant access

This chapter includes the following topics:

Prerequisites of instant access

If you are using instant access, ensure that the WORM instance can access the following port on vCenter:

Table 12-1 Port details

Instance	VMware component	Port number
WORM	vCenter	443

Things to consider before you use the instant access feature

Note the following about the **Instant access virtual machines** feature:

- This feature is supported with backup copies that are created from the local or cloud LSU (logical storage unit) using the NetBackup web UI or Instant Access APIs.
  For more information about limitations of instant access for cloud LSU (logical storage unit), refer to the `NetBackup Deduplication Guide`.
- This feature is supported with backup copies that are created from protection plans or policies.
- This feature is supported for NetBackup Appliance, NetBackup Virtual Appliance, Flex Appliance, and Build Your Own (BYO) server.
  Instant access on Flex WORM storage requires the following services:
  - NGINX, NFS. SAMBA, WINBIND (if Active directory is required), SPWS, VPFS
- This feature is limited to 50 concurrent mount points from a Media Server Deduplication Pool (MSDP) media server or from a WORM storage server. If you have a Flex appliance, this feature is limited to 50 concurrent mount points from each node.
- By default, vSphere allows a maximum of eight NFS mounts per ESXi server. Note that NetBackup requires an NFS mount for each instant access VM you create. To remove the NFS mount, remove the instant access VM when you are done with it.
  If the NFS limit for an ESXi host has been reached and you try to create another instant access VM, the attempt fails. To increase the maximum NFS mounts per ESXi server, see the following VMware article:
  https://kb.vmware.com/s/article/2239
- This feature does not support backups of VMs that have independent disks. VMware does not support snapshots of independent disks in a VM, either persistent disks or non-persistent disks. As a result, independent disks are not backed up.
  For more information on independent disks and NetBackup, see the following article:
  https://www.veritas.com/docs/000081966
- This feature does not support VMs that have the disks that were excluded from the backup. For a policy, on the Exclude Disks tab select No disks excluded. For a protection plan, clear the Exclude selected virtual disks from backups check box.
- This feature does not support VMs that have a disk in raw device mapping mode (RDM) or that have a disk in Persistent mode.
- For Windows restore, the ReFS file system is not supported.

- The version of the ESXi server that is used to create a VM using **Instant access virtual machines** must be equal to or newer than the version of the ESXi server that contains the VM backup images.
- For file or folder download with the **Download** option, the NetBackup web UI must be able to access the media server with the same name or IP address that the primary server uses to connect to that media server.
  See Troubleshooting VMware backups on page 332.
- If the media server appliance uses a third-party certificate, you need to create certain configurations on the NetBackup primary server before you use this feature.
  For more information, refer to the "Third-party certificates" and "Implementing third-party SSL certificates" sections in the NetBackup Appliance Security Guide.
- This feature does not support restore of multiple files or folders, which are located in different volumes, partitions, or disks.
- Use the Windows administrator account credentials when you restore multiple files or folders to a Windows VM. You must be logged on to the target Windows VM with these account credentials.
- Some ACL entries are not in the restored file because ACL entries for these users or groups cannot be restored. For example, TrustedInstallers, All Application Packages.
- The Instant Access feature does not support a Windows 10 compact operating system. To verify if your operating system is compressed, run `compact "/compactos:query"` on the command prompt before backing up your VM.
  To disable the compression, run `"compact /compactos:never"` on the command prompt before backing up your VM. You can then use the Instant Access feature for your VM backups.
- To restore files and folders, the target VM must be in a normal state, and not in a sleep or hibernate mode.
- A 5-minutes-alive-session threshold is defined in Appliance and BYO web server NGINX. The files and folders that are selected for download must be compressed and downloaded within this threshold.
- To create an instant access virtual machine, you must have read and write access to the VMware data center where the virtual machine is created.
- To ensure that Instant Access works effectively after the storage server and primary server are upgraded from an earlier NetBackup version, restart the NetBackup Web Service on the upgraded primary server with the following commands:
  `/usr/openv/netbackup/bin/nbwmc stop`

```
/usr/openv/netbackup/bin/nbwmc start
```

- If you have to download or restore files or folders from a Windows VM, ensure that the number of Windows registry hives are less than 10000.
  More information is available about registry hives.
- An image cannot be deleted if an instant access VM is created from it. The instant access feature uses data from a backup image. If the image is expired, the data might be unavailable and the instant access VM may face data loss. After the instance access VM is deleted, the image can be expired.
- The instant access feature does not support hard links. If you create a universal share from an image and the image has hard link files, `vpfsd` shows show these hard link files as having 0 bytes size.
- Instant access supports the DataSets feature from vSphere 8.0.

Create an instant access VM

You can create an instant access VM from a NetBackup backup image. The VM is available almost instantaneously, achieving a near-zero recovery time objective. NetBackup mounts the VM's snapshot directly on the backup storage device to allow your ESXi host or cluster to treat the snapshot as a normal VM.

The mounted VM snapshot can be used for a variety of purposes. For example:

- Recovering files from the VM, or copying a vmdk file.
- Running tests on the VM, such as testing a patch.
- Troubleshooting or disaster recovery.
- Verifying an application.

Note: This feature is supported for NetBackup Appliance, NetBackup Virtual Appliance, Flex Appliance, and Build Your Own (BYO) server. This feature requires that the NetBackup backup image is stored on a Media Server Deduplication Pool (MSDP) storage device. More information on using instance access VMs is available:

See “Things to consider before you use the instant access feature” on page 182.

To create an instant access VM

1. On the left, click **VMware**.
2. Locate the VM and click on it.

3 Click the **Recovery points** tab, then click the date on which the backup occurred.

The available images appear in rows with the backup timestamp for each image.

4 Select the **Allow the selection of recovery of points that are malware-affected** to be able to recover malware affected images. This option is only available for the recovery points which contains malware-affected images.

Note: This options is enabled only for users with required permissions.

---

5 On the image or the copy of the image that has the option to recover using instant access, click **Recover > Create instant access virtual machine**.

6 Review the recovery settings and make changes if needed.

Note the **Recovery options**:

Allow overwrite of existing virtual machine	If a VM with the same display name exists at the destination, that VM must be deleted before the recovery begins. Otherwise, the recovery fails.
Power on after provisioning	Automatically powers on the VM when the recovery is complete.
Enable vMotion	Starts the migration of the VM after it is created and then displays progress of the VM migration. **Note:** For a NetBackup 8.1.2 storage server, the vMotion option is not used even if it is enabled.

7 Click **Create**.

NetBackup makes a snapshot of the VM backup image and creates an instant access mount point. The snapshot of the image appears on the **Instant access virtual machines** tab. You can now use the VM like any other VM on the ESXi server.

8 For details on the restored VM, click on the VM under the **Instant access virtual machines** tab and click **View details**.

9 When you are finished with the VM, you can click **Delete** to remove the mounted VM snapshot. The VM is removed from the ESXi server.

Note: If vMotion is enabled and completed successfully, deleting a VM only removes the mounted share. The VM is still available on the ESXi server as this VM is migrated to another datastore.

Restore files and folders from a VM backup image

You can browse an instant access image of the VM to restore files and folders.

Note: More information on using instance access VMs is available:

See "Things to consider before you use the instant access feature" on page 182.

To restore files and folders from a VM backup image

1 On the left, click **VMware**.

2 Locate and click on the VM.

3 Click the **Recovery points** tab. In the calendar view, click the date on which the backup occurred.

The available images are listed in rows with the backup timestamp for each image.

4 Select the **Allow the selection of recovery of points that are malware-affected** to be able to recover malware affected images. This option is only available for the recovery points which contains malware-affected images.

Note: This option is enabled only for users with the necessary RBAC role or related RBAC permissions.

5 On the image or the copy of the image that has the option to recover using instant access, click **Recover > Restore files and folders**.

NetBackup creates an instant access mount point in the background.

6 Select the files and click **Add to restore list**.

Click on a folder to drill into it. Use the folder path to navigate back to higher levels in the hierarchy.

Enter a file name to search for files.

The restore list displays the selected files and folders with the location and the estimated size of each file.

7 Select the restore options:

- **Restore everything to the original directory**
  You can manually enter the credentials or select an existing credential:
  - **Manually enter credentials**: Enter the name of the target VM (the default is the original VM) and the username and password for the target VM.
    Or
  - **Select existing credentials**: From the list of available credentials, select the credential and click **Next**.
- **Restore everything to a different directory**
  - In **Directory for restore**, enter the destination path for restore.
  - Select the **Flatten existing directory structure** check box to restore all files to a single directory.
  - You can manually enter the credentials or select an existing credential:
    - **Manually enter credentials**: Enter the name of the target VM (the default is the original VM) and the username and password for the target VM.
      Or
    - **Select existing credentials**: From the list of available credentials, select the credential and click **Next**.

8 Select the required check box:

- **Append string to file names**: Append the specified string to the destination file names before any file extension. This value only applies to files.
  Overwrite existing files: Overwrites all the existing files.
- **Restore directories without crossing mount points**: Skips over file systems mounted in the selected directories.

- **Guest VM uses Windows User Account Control (UAC)**
- **Create new files for hard links**: Creates a new file that is associated with the hard link.
- **Rename targets for soft links**: Creates a link that references the new target.
- **Allow recovery of files infected by malware**: By default infected files are not recovered. This allows the user to change the default behavior of clean recovery.

Note: This option is displayed only when you select the **Allow the selection of recovery of points that are malware-affected** option.

A summary of your selections is displayed.

9 Click **Start recovery** to restore the files.

The **Activity** tab displays the status of the recovery.

Download files and folders from a VM backup image

You can browse an instant access image of the VM to download files and folders.

Note: More information on using instance access VMs is available:

See "Things to consider before you use the instant access feature" on page 182.

To download files and folders from a VM backup image

1 On the left, click **VMware**.

2 Locate and click on the VM.

3 Click the **Recovery points** tab. In the calendar view, click the date on which the backup occurred.

The available images are listed in rows with the backup timestamp for each image.

4 Select the **Allow the selection of recovery of points that are malware-affected** to be able to recover malware affected images. This option is only available for the recovery points which contains malware-affected images.

Note: This options is enabled only for users with required permissions.

5 On the image or the copy of the image that has the option to recover using instant access, click **Recover > Download files and folders**.

6 Select the files and click **Add to download list**.

Click on a folder to drill into it. Use the folder path to navigate back to higher levels in the hierarchy.

Enter a file name to search for files.

The download list displays the selected files and folders with the location and estimated size of each file.

7 After the download package is created, click **Download**.

The **Activity** tab displays the status of the recovery.

Instant access Build Your Own (BYO)

You can build your own VMs (with Red Hat enterprise operating system) to support VMware instant access. You can use the following features:

- Create instant access VMs.
- VMware vMotion.
- Download files and folders.
- Restore files and folders.

To use instant access with a BYO VM created with an earlier NetBackup release, you must upgrade to NetBackup 8.3.

Prerequisites of Instant Access Build Your Own (BYO)

Prerequisites (fresh install and upgrade):

- The BYO storage server with Red Hat Enterprise Linux 7.6 and later, same as the NetBackup Appliance operating system version.
- The BYO storage server with docker/podman installed.
  - The docker/podman version must be same as the one in the corresponding official RHEL version release. You need to install it from the corresponding RHEL yum source (RHEL extra).
  - The docker/podman application is included in the environment path.
- The BYO storage server with NFS service installed.
- The BYO storage server with NGINX version installed.
  - The NGINX version must be same as the one in the corresponding official RHEL version release. You need to install it from the corresponding RHEL yum source (epel).
  - Ensure that the `policycoreutils` and `policycoreutils-python` packages are installed from the same RHEL yum source (RHEL server) and then run the following commands:
    - `semanage port -a -t http_port_t -p tcp 10087`
    - `setsebool -P httpd_can_network_connect 1`
  - Ensure that the `/mnt` folder on the storage server is not mounted by any mount points directly. Mount points should be mounted to its subfolders.
  - Enable the logrotate permission in selinux using the following command:
    `semanage permissive -a logrotate_t`
- For BYO, docker/podman container is used to browse VMDK files. Data related to the container is stored at the following location: `/var/lib/` and requires minimum 20 GB free space.

Hardware configuration requirement of Instant Access Build Your Own (BYO)

Table 12-2 Hardware configuration requirement

CPU	Memory	Disk
■ Minimum 2.2-GHz clock rate. ■ 64-bit processor. ■ Minimum 4 cores; 8 cores recommended. For 64 TBs of storage, the Intel x86-64 architecture requires eight cores. ■ Enable the VT-X option in the CPU configuration.	■ 16 GB (For 8 TBs to 32 TBs of storage - 1GB RAM for 1TB of storage). ■ 32 GBs of RAM for more than 32 TBs storage. ■ An additional 500MB of RAM for each live mount.	Disk size depends on the size of your backup. Refer to the hardware requirements for NetBackup and Media Server Deduplication Pool (MSDP).

Frequently asked questions

Here are some frequently asked questions for instant access Build Your Own (BYO).

Table 12-3 Frequently asked questions

Frequently asked question	Answer
How can I enable instant access file browsing (for file download and restore) on BYO after the storage is configured or upgraded without the docker/podman installed?	Perform the steps in the following order: **1** Install the required docker/podman version. **2** Start using the Instant Access feature. For example, you can download files, restore files, and so on.
How can I enable the VMware instant access feature on BYO after storage is configured or upgraded without the nginx service installed?	Perform the steps in the following order: **1** Install the required nginx service version. **2** Ensure that the new BYO nginx configuration entry: `/etc/nginx/conf.d/byo.conf` is part of the HTTP section of the original: `/etc/nginx/nginx.conf` file. **3** Run the command: `/usr/openv/pdde/vpfs/bin/vpfs_config.sh --configure_byo`

Table 12-3 Frequently asked questions *(continued)*

Frequently asked question	Answer
How can I resolve the following issue in the vpfs-config.log file that is raised from: `Verifying that the MSDP REST API is available via https on port 10087`	Perform the steps in the following order: **1** Install the `policycoreutils` and `policycoreutils-python` packages through yum tool. **2** Add the following rules that SELinux requires for Nginx to bind on the 10087 port. ■ `semanage port -a -t http_port_t -p tcp 10087` ■ `setsebool -P httpd_can_network_connect 1` **3** Run the following command: `/usr/openv/pdde/vpfs/bin/vpfs_config.sh --configure_byo`
Instant Access for BYO uses a self-signed certificate by default and only supports *.pem external certificate. How do I replace it with a certificate signed by external CA (*.pem certificate), if required?	To configure the external certificate, perform the following steps. If the new certificate is already generated (the certificate must contain long and short host names for the media server), go to step 4. **1** Create the RSA public or private key pair. **2** Create a certificate signing request (CSR). The certificate must contain long and short host names for the media server. **3** The External Certificate Authority creates the certificate. **4** Replace `<PDDE Storage Path>/spws/var/keys/spws.cert` with the certificate and replace `<PDDE Storage Path>/spws/var/keys/spws.key` with the private key. **5** Run the following command to reload the certificate: `/usr/openv/pdde/vpfs/bin/vpfs_config.sh --configure_byo`

Table 12-3 Frequently asked questions *(continued)*

Frequently asked question	Answer
How can I disable media automount for the instant access livemount share in gnome? If the automount is enabled, the source folder is mounted from the livemount share in gnome and smaller disks appear. In this scenario, the instant access feature does not work properly. The mounted disk content source is from the `.../meta_bdev_dir/...` folder under livemount share, while the mount target is in the `/run/media/...` folder.	Follow the guideline to disable the gnome automount: https://access.redhat.com/solutions/20107

How can I resolve the following issue in the `/var/log/vpfs/vpfs-config.log` file?
 	
`**** Asking the NetBackup Webservice to trust the MSDP webserver (spws) **** /usr/openv/netbackup/bin/nblibcurlcmd failed (1):`
	Perform the steps in the following order:
	
	1 Ensure that your NetBackup primary server is up and there is no firewall blocking the connection between the NetBackup primary server and storage server.
	
	2 Run the following command on storage server to verify the connection status:
	
	`/usr/openv/netbackup/bin/bpclntcmd -pn`
	
	3 After the NetBackup primary server is up and connection between the NetBackup primary server and storage server is allowed, run the following command:
	
	`/usr/openv/pdde/vpfs/bin/vpfs_config.sh --configure_byo`

Why can't I use instant access to browse, download, or restore files from the VMware backup image?	For security reasons, the directory `/usr/openv/pdde/pdopensource/supermin_appliance` has been removed from `VRTSpddes.rpm` package in version 10.5.0.1 and later versions of the BYO storage server. As a result, VMs with the BTRFS file system or using Windows Logical Disk Manager (LDM) experience the following: ■ Cannot use instant access to browse, download, or restore files from the VMware backup image. ■ Malware scanning does not work for these VMs. To restore support for these features, you must manually download and install certain packages. For detailed instructions, see the following article: https://www.veritas.com/content/support/en_US/article.100071923

VM malware scan

You can create a malware scan livemount from a NetBackup image with NetBackup Recovery API. The livemount exports all VM files and folders via NFS or SMB protocol instantly, which allows NFS or SMB client mount the export path and do malware scan on the exported VM files and folders.

This feature provides the following malware scan APIs:

- POST /recovery/workloads/vmware/malware-scan-mounts
- GET /recovery/workloads/vmware/malware-scan-mounts
- GET /recovery/workloads/vmware/malware-scan-mounts/{mountId}
- DELETE /recovery/workloads/vmware/malware-scan-mounts/{mountId}.

For more details, refer NetBackup 10.0.1 API Reference on SORT

Chapter 13

Instant rollback

This chapter includes the following topics:

Prerequisites of instant rollback

The prerequisites for Instant Access Build Your Own (BYO) are also applicable to the Instant Rollback feature.

See “Prerequisites of Instant Access Build Your Own (BYO)” on page 191.

For NetBackup FlexScale, the software packages that instant rollback requires are included with the NetBackup FlexScale deployment. For more information, see the *NetBackup Flex Scale Administrator's Guide*.

If you are using instant rollback, ensure that the WORM instance can access the following ports on vCenter and ESXi servers:

Table 13-1 Port details

Instance	VMware component	Port number
WORM	vCenter	443
WORM	ESXi host(s)	902

Things to consider before you use the instant rollback feature

Note the following about the instant roll back virtual machines feature:

- This feature is supported with backup copies. These copies are created with protection plans or classic policies.
- This feature is supported for NetBackup Appliance, NetBackup Virtual Appliance, Build Your Own (BYO) server, and NetBackup FlexScale.
- This feature does not support backups of VMs that have independent disks. VMware does not support snapshots of independent disks in a VM, either persistent disks or non-persistent disks. As a result, independent disks are not backed up.

  For more information, see the following:

  https://www.veritas.com/docs/000081966
- This feature does not support VMs that have the disks that were excluded from the backup. For a policy, on the **Exclude disks** tab select **No disks excluded**. For a protection plan, clear the **Exclude selected virtual disks from backups** checkbox.
- This feature does not support VMs that have a disk in raw device-mapping mode (RDM).
- This feature lets you select a maximum of 100 VMs for rollback at a time. If you select more than 100 VMs the **Roll back instantly** option is not displayed.

  For example, if you want to rollback 180 VMs, you need create two rollback requests for the same job. One for 100 VMs and the second for 80 VMs.
- In this feature, one instant rollback VM requires one livemount. Each livemount can be retained for one day. So the number of VMs that can support roll back depend on the total number of livemounts available. By default, the livemounts value is set to 200.

  You can change this default value from the following location: `storage path/spws/etc/spws.cfg`

  MaxAllowedLivemounts=200

  For NetBackup FlexScale, the livemounts value is set to 100 by default on each MSDP engine in MSDP cluster.

  You can change this default value from the following location for MSDP engine: `/msdp/data/dp1/pdvol/spws/etc/spws.cfg`

Note: The total livemount number configured in instant rollback, VMware instant access, MSSQL instant access, and universal share must not exceed the **MaxAllowedLivemounts** value.

- This feature does not support the add, remove, or update DataSets feature for virtual machines. The Instant rollback feature does not roll back DataSets.

Instant rollback from a VM backup image

NetBackup 9.1 and later lets you roll back a VM instantly from a backup image. Only backup images that support instant access can support instant rollback.

You can perform instant rollback for multiple VMs. You can also roll back a VM multiple times to any recovery point.

For example, if you have three backup images, B1, B2, and B3, you can first roll back the VM to B1, then to B3, then to B2, and so on.

After the rollback is completed, all data after the selected recovery point is no longer available.

To instantly roll back from a VM backup image

1. On the left, click **VMware**.
2. To select the backup image, do one of the following:

Click the VM	1. Locate the VM and click on it. 2. Click the **Recovery points** tab, then click the date on which the backup occurred. The available images appear in rows with the backup timestamp for each image. 3. Select the **Allow the selection of recovery of points that are malware-affected** to be able to recover malware affected images. This option is only available for the recovery points which contains malware-affected images. **Note:** This options is enabled only for users with required permissions. 4. On the image or a copy of the image, click **Recover > Roll back instantly**.

Select the check box

1. Select the check box corresponding to the VM that you want to roll back and click **Roll back instantly**.

   You can select multiple VMs to perform instant rollback.

2. Select any one of the roll back options:
   - **Roll back to: Most recent**
     NetBackup displays the most recent instant access recovery points in a month.
   - **Roll back to: Before specific date and time**
     Select the date and time.
     NetBackup displays the most recent instant access recovery points going a month before the selected date and time.

     Note: NetBackup displays a warning about malware affected images.
   - Select the **Allow the selection of recovery of points that are malware-affected** to be able to recover malware affected images. This option is only available for the recovery points which contains malware-affected images.

   Note: This options is enabled only for users with required permissions.

3. Click **Roll back**.

Use the **Actions** menu

1. Click **Actions > Roll back instantly** corresponding to the VM that you want to roll back.
2. Select any one of the roll back options:
   - **Roll back to: Most recent**
     NetBackup displays the most recent instant access recovery points in a month.
   - **Roll back to: Before specific date and time**
     Select the date and time.
     NetBackup displays the most recent instant access recovery points going a month before the selected date and time.
   - Select the **Allow the selection of recovery of points that are malware-affected** to be able to recover malware affected images. This option is only available for the recovery points which contains malware-affected images.
3. Click **Roll back**.

3. Select the wanted options and then click **Roll back**.

   The **Activity monitor** tab displays the status of the rollback.

Chapter 14

Continuous data protection

This chapter includes the following topics:

About continuous data protection

Continuous Data Protection (CDP) is a smart way to capture fast copies of backups for the VMware VMs, without stunning the VMs. Using CDP, you can rapidly make recent copies of backups and use NetBackup to retain and restore the backups as required.

Here are some salient features of CDP:

- Completely web UI-based protection and recovery of VMware VMs.
- Versatile API-based protection.
- You can use traditional VADP-based backups along with CDP for VMware. The backup images are independent of each other, and they are treated separately for incremental backup or recovery purpose.
- Bring Your Own Device (BYOD): You can use a Red Hat Linux-based NetBackup media server as a CDP gateway.
- Support for ESXi and various datastore types. Refer to the Software compatibility list for the latest information.
- Accelerator-based backup. Support for accelerator enabled storage like MSDP and OST.
- Support for Instant access. You can start the VMs from MSDP storage.
- Agentless single file restore from MSDP.
- RBAC support for entire protection and restore workflow.
- Traditional and capacity-based licensing.
- The Veritas Resiliency Platform is fully compatible with the Veritas I/O filters used by CDP.

CDP terminology

The following table describes the concepts and terms that are used in Continuous Data Protection (CDP).

Table 14-1 CDP terminology

Term	Explanation
CDP gateway	CDP configured media server.

Table 14-1 CDP terminology *(continued)*

Term	Explanation
VAIO	The VMware framework consists of vSphere APIs for I/O filtering. This framework lets CDP run filters on ESXi servers and intercept I/O requests from guest operating systems to a virtual disk.
Full sync	NetBackup retrieves a VM's entire data from the ESXi.
OST	Open Storage Technology is a STU supported by NetBackup.
MSDP	Media server deduplication storage pool is a NetBackup dedupe technology engine to optimize backup storage.
Storage policy	A VMware vSphere feature that enables storage profile creation by administrators. This means that the VMs do not need to be configured separately, and management may be automated.
VIB	vSphere Installation Bundle. At a conceptual level, a VIB is similar to a tarball or compressed archive. It is a collection of files packaged into a single archive to facilitate distribution.
nbcctd	CDP service (daemon) running on the CDP gateway.
Staging area	A storage location on the CDP gateway where NetBackup temporarily stores I/Os received from the ESXi.
Storage quota	Allocated limited storage size for VMs using CDP protection.
Reserved quota	Shared storage between all VMs registered to a CDP gateway.
VADP	VMware VADP is a VMware vStorage API that backs up and restores vSphere virtual machines (VMs).

CDP architecture

The CDP gateway is configured on a NetBackup media server. Once the configuration is done, NetBackup starts the `nbcctd` daemon on the CDP gateway. This process services all I/Os from ESX and enables other NetBackup components on the gateway to take backup. To back up this data, you also need to configure an MSDP or OST accelerator-based STU. You can configure multiple CDP gateways and MSDP/OST accelerator-based STUs as required. NetBackup REST APIs for CDP leverage this feature. Refer to NetBackup REST APIs Swagger documentation for more information.

Figure 14-1 CDP architecture

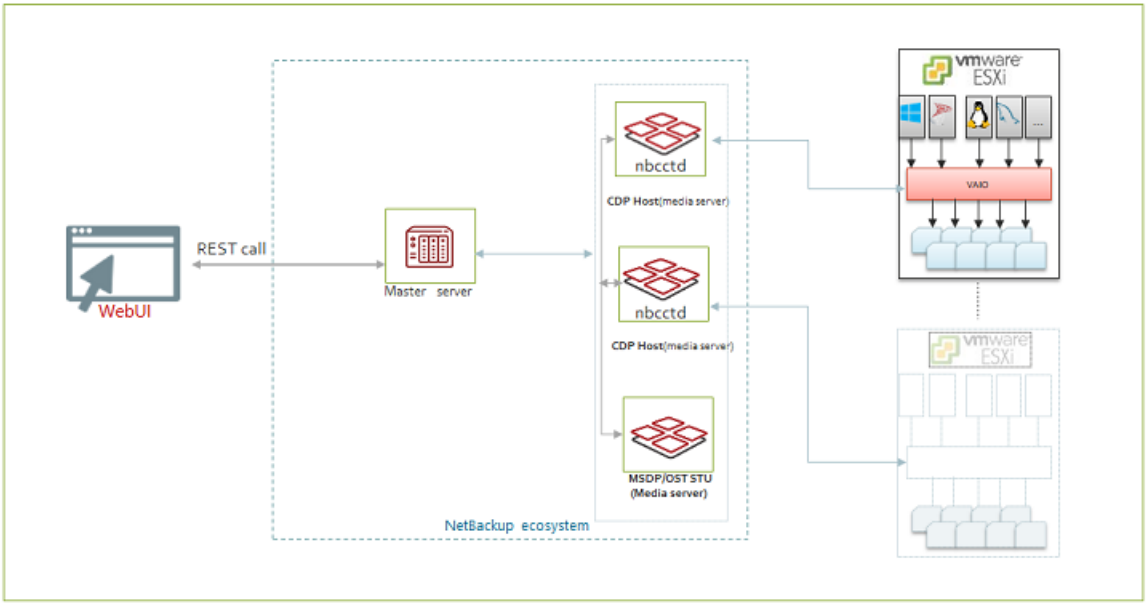

Prerequisites

Prerequisites for using CDP

- CDP for VMware exclusively supports accelerator-based backup. So, CDP needs accelerator-compliant storage units based on MSDP or OST-based storage.
- CDP uses a file system as a staging area on the CDP gateway. See the Software compatibility list for the supported file systems.
- The media server that is associated with MSDP should have NetBackup version 9.1 or higher.
- Capacity-based and traditional license for enabling the feature.
- The port 33056 on the CDP gateway must be open for the ESXi server to communicate to the CDP gateway.
- VMware server credentials need privileges for NetBackup to start, stop, restart, and refresh the Common Information Model (CIM) service on the ESXi host.
- You can configure a CDP gateway on the RHEL-based NetBackup media server platform.

- Create a VMware storage policy for replication using the VAIO component. Attach the storage policy to each disk of the VMs that you want to protect using CDP. For details, see the Veritas Technical Support knowledge base article on How to create vtstap storage policy in VMware vCenter.

Veritas IO filter for VAIO requirement

You can download and deploy the VAIO drivers package, version 4.0.0, to use with your CDP deployment. Refer to the Software compatibility list for the latest version and information on how to download it.

You must install the vSphere Installation Bundle (VIB) on the vCenter cluster before configuring protection in NetBackup. Note that you do not need to deploy VIB on vCenter for restore purposes. See the Veritas Technical Support knowledge base article on *Deploying an IO Filter solution to a cluster using VMware MOB.*

Storage Policy requirements

Before you can deploy CDP, you need to create a VM storage policy. The storage policy must have a component chosen as "Replication" and a provider as "vtstap". This policy must be attached to each disk of the VM to be protected. Otherwise, backup jobs fail. For details, see the Veritas Support knowledge base article on How to create vtstap storage policy in VMware vCenter

---

Note: Detaching the storage policy results in loss of protection for the VM. If you detach the Veritas IO filter storage policy from a VM, I/O tapping for the VM is stopped, so the data from this VM does not get saved in the CDP gateway. Hence, the consequent backup jobs remain blank backup jobs, even after all the data from the CDP staging area is moved to the backup storage. So, we recommend removing protection of the VM(s) from the NetBackup protection plan once you detach the vtstap policy from the VM(s).

Capacity-based licensing for CDP

Licensing collects the total number of front-end terabytes protected by NetBackup. The front-end data size for a CDP backup is nearly the same as the consumed storage size on the ESX datastore by the VMs.

The `nbdeployutil` utility reports data usage for the VMs. These rules are applied to report data size:

- Calculate the total number of bytes written during backup (X) and the VM size from the ESX datastore (Y). The reported size is the smaller value of X and Y.
- If different policies use the same virtual machine, the policy with higher data size is accounted.

- If a VADP and CDP policies protect the same VM, then you are charged only once, with the higher size.

Administrators can use the following steps to verify the data size reported by licensing:

- Verify the size of the VMs on the ESX datastore on the vCenter. Navigate to **Datastore** > **Files** > **VM**, the **Size** column shows the size occupied on datastore.
- Verify the bytes written during backup for the same VM.
- Calculate the minimum of the above two values.

Steps to configure CDP

To configure CDP for your workload, you must perform the following tasks.

Operations on the VMware vCenter

1. Install the I/O filters by Veritas. See Veritas Support knowledge base article on Deploying an IO Filter solution to a cluster using VMware MOB.
2. Attach the storage policy to ESXi. For details, see Veritas Support knowledge base article on How to create vtstap storage policy in VMware vCenter

Operations on the NetBackup console

1. Create an MSDP or OST-based storage for the backup destination. See the information on how to configure storage in the *NetBackup Web UI Administrator's Guide*.
2. Create a CDP gateway.
3. Create a CDP-based protection plan for your VMware workload. See the *Managing protection plans* chapter of the *NetBackup Web UI Administrator's Guide*.
4. Protect the required VMs with the protection plan.
5. Monitor jobs.

Removing VMs from the CDP gateway

When CDP protection is no longer required for a VM, you can remove protections from that VM, or switch the VM to classic policy.

To remove CDP protection from a VM

1. Go to vCenter and change the VM's storage policy from `vtstap` to `Datastore default`.
2. In the NetBackup web UI, on the left, click **VMware** under **Workloads**. You can see a list of VMs with protection details.
3. Click the name of the VM, from which you want to remove protection, in the subsequent page, click **Remove protection**.

   You can see a confirmation message when the VM is removed.

If you remove protection from a VM without removing the `vtstap` policy from the VM, you can see a partially successful removal message in the UI. These partially removed VMs are not included in the **Total VMs subscribed** count in the **Continuous data protection gateway** tab.

Note: The partially removed VMs are neither protected by CPD nor by any classic policy. Also, you cannot re-subscribe the VM to a CDP gateway. Hence, it is recommended to detach the `vtstap` storage policy from the VM, and fully unsubscribe the VM from the CDP gateway.

Defining the CDP gateway

You need to define a gateway for your CDP deployment, before you can protect any VMs. You can define the CDP gateway in a VM that is a NetBackup media or primary server.

Note: Before defining the CDP gateway ensure that your system time is synchronized with the network time.

To define a CDP gateway

1. On the left, click **VMware** under **Workloads**.
2. On the top right, click **VMware settings**, and click **Continuous data protection gateway**.
3. Click **Add**. Enter a **Host name** and **Storage path**. The storage path should have an independent file system, other than the root. Do not share this same location with other applications like, MSDP.

4 On the next page, if your gateway version is 9.1. specify the parameter **Maximum number of concurrent jobs**, as described in the subsequent table, and click **Save** to save the gateway.

If your gateway version is 10.0, click **Advanced** to specify the advanced parameters to configure and fine-tune your CDP gateway. You can also use this set of parameters to estimate how many VMs you can support using CDP protection for a particular configuration of the gateway.

Parameter	Description
Maximum number of concurrent jobs.	The maximum number of CDP jobs that can run simultaneously in the gateway. A higher number may indicate increased peak resource consumption.
Maximum number of simultaneous initial sync	Number of VMs that can take full backup simultaneously during the initial phase of CDP protection. Specifying a higher value than the default, may cause increased resource consumption and affect existing protection.
Reserved memory for Continuous data protection	Reserved memory for the gateway. Enter a value in GB that is equal to or smaller than 90% of the total physical memory.
Data staging area per VM	Specify storage for each VM.
Reserved staging area	Additional storage area to handle the I/O spikes in the VMs.

5 Click **Estimate the number of VMs** to calculate how many VMs this gateway can support for this given configuration.

6 Click **Save**, to add the gateway.

Sizing considerations

This section describes the CDP gateway's sizing requirements, based on your environment's workload.

Note: If you plan to support a large number of VMs using the CDP gateway, deploy the CDP gateway and the MSDP or media server hosting the storage unit, on different hosts.

Note: If the CDP gateway and MSDP are co-located on the same media server, then the CDP service consumes 20% of available memory (RAM) for its internal use. If the CDP gateway is standalone on the media server, it consumes 50% of the available memory for the same. From NetBackup version 10.0 onwards, you can configure this value in the UI.

Gateway sizing

You need to size the CDP based on the number of VMs that you want to protect. Consider the requirements described in this section, while calculating the requirements for the gateway.

CDP enables you to continuously tap the I/Os done by the VMs. NetBackup, by default, uses 10-GB storage space on the staging area per VM. When IO tapping starts, the CDP service starts writing the data into this 10-GB storage. Once this storage limit is reached, the CDP service (nbcctd) initiates a backup job to move this data from the gateway to the backup storage.

Out of the total available space on the CDP staging path, by default, NetBackup reserves 25% for usage beyond the allocated storage per VM. This storage is common for the subscribed VMs to the gateway. See "Defining the CDP gateway " on page 206. , for how to do it on version 10.0 onwards. You can reconfigure this value in the `nbcct.conf` file in NetBackup 9.1.

To configure reserved storage in NetBackup 9.1

1. Log on to the CDP gateway.
2. Navigate to the `<stagingloc>/nbcct/` directory, and open the `nbcct.conf` file in a text editor.
3. Enter the required values against the parameters *CCT_VM_QUOTA_SIZE_IN_MB* and *CCT_VM_QUOTA_RESERVE_PERCENT*
4. Restart the `nbcctd` service.

Storage requirement for the gateway

When NetBackup receives the data from the ESXi IO daemon, it stores the data in the in-memory cache. Recommended is a minimum of 160 MB of data for each VM.

For example, you protect 40 VMs in a gateway. So, you need 40*160 MB = 6400-MB RAM. Allocating more RAM increases the in-memory cache size when the CDP service starts, ultimately increasing the IO performance of the service.

Similarly, to stage 40 * 10-GB = 400-GB (75%) + 134GB (25%) reserved, that is approximately 540-GB space you need to have in the staging area.

Increasing per-VM storage allows to NetBackup to back up more data per backup job. Increasing the reserved storage for the CDP gateway lets you receive more data without any interruption to the protection. Note that even when the staging path is fully occupied, it does not affect the applications inside the VM. NetBackup catches up the data produced by applications during that time, and moves it to the backup storage in the subsequent backup jobs.

Note: If NFS is used for the staging area, the minimum required throughput is 100 MB/sec.

First 24-hours experience

When you start using the CDP feature, it is important to observe the system and tune it according to your business demands. Add any required hardware configurations to maximize protection and performance. First, you can use default values and start subscribing the VMs according to the requirements mentioned in this section. You should check the following:

- Number of immediate backup jobs that the CDP service initiates due to the staging storage in full condition.
- You can check the CDP backup engine notifications on NetBackup web UI.
- Underlying provisioned storage performance. Like the NetBackup installation disk, CDP staging area, and MSDP storage disks.
- Network utilization and available bandwidth.
- CPU and memory consumption when receiving data from the ESXi, and when the backup jobs are running.

Note: If you observe slow I/Os from the I/O daemon, check network bandwidth and system RAM. See “Defining the CDP gateway ” on page 206. , for how to increase the in-memory cache size in NetBackup 10.0 onwards. For NetBackup 9.1, you can do it using the *CCT_POOL_SIZE_QUOTA_PERCENTAGE* parameter in the `nbcct.conf` file.

Limiting concurrent CDP backup jobs

You can set a limit for the simultaneous CDP snapshot jobs that can run on the CDP gateway at a time. For example, if you protect 20 VMs and you set a limit of 5, then only 5 VMs can run simultaneous backups, and 15 VMs stay in queue. This setting is required for optimized use of your system and network resources. By default, the resource limit value is 0, representing no limit.

See “Defining the CDP gateway ” on page 206. for information on how to do it on NetBackup version 10.0 onwards. For NetBackup 9.1 follow the procedure described below.

To set a resource limit, we have the following API:

```
POST /config/resource-limits

{
  "data": [
    {
      "type": "resource-limits",
      "id": "string",
      "attributes": {
        "resources": [
          {
            "resourceType": "string",
            "resourceName": "string",
            "resourceLimit": 0,
            "additionalData": "string"
          }
        ]
      }
    }
  ]
}
```

Here,

- `Id` represents the workload, which is `Cdp`
- `resourceType` should be `Cdp-Backup`
- `resourceName` represents the CDP gateway host name. It should be the same as specified in the protection plan. If you keep an empty string for `resourceName`, the `resourceLimit` value is set as a global limit, which is applicable to all the configured CDP gateways.
- The `resourceLimit` value sets the value of backup jobs for that gateway.

To retrieve the list of resource limits for a CDP workload type, use:

```
GET - /config/resource-limits/cdp
```

To update the value of `resourceLimit` for particular gateway, hit the POST API with the change in `resourceLimit` for the same record.

To delete the specified granular resource limits, use:

```
DELETE - /config/resource-limits
```

Only the resource limit set for a particular resource can be deleted. Provide both the resource type and the specific resource of that type.

Controlling full sync

When you subscribe a VM to a CDP-enabled protection plan, NetBackup initiates full sync, to get the entire data of the newly protected VM. For a newly subscribed VM, NetBackup does not have any data to apply the incremental backup features; hence full sync is initiated. During a full sync, NetBackup captures the entire data of the VM, from the underlying VMDKs to the CDP staging location, and subsequently to the NetBackup STUs.

Full sync is normally triggered when you subscribe a new VM to a CDP-enabled protection plan, but in certain scenarios, you can manually initiate a full sync:

- Accidental corruption or deletion: CDP maintains backed up data of the VMs at the staging location in proprietary format files. If these files for a VM are accidentally deleted or corrupted, the subsequent backup job for the VM fails, citing data integrity mismatch. In this case, you can initiate a force rescan schedule backup, and subsequently, a full sync of the VM takes place.
- Following a manually triggered force-rescan schedule.
- CDP service can initiate full sync to receive VM data whenever necessary.

During full sync, data flows from the ESXi to the CDP gateway. Depending on the data size of the VMs, the volume of this data can be substantially large that can consume plenty of resources like network, memory, processing power, and storage. This also affects the backup operations of the VMs subscribed earlier.

If you subscribe more than 5 VMs at a time, say 7, then, full sync is initiated for 5 VMs, and 2 are in wait state.

Therefore, it is recommended to limit the number of concurrent full sync operations to optimize system resources. The default number of concurrent full sync is 5. This allows 5 VMs to perform full sync concurrently. Other VMs needing full sync need to wait in a queue. This way, the system resources are managed optimally.

Recommendation for controlling full sync:

- Subscribe the VMs in batches of five or less.
- Once a subscribed VM completes full sync, you can see a message in the UI, then you can proceed to subscribe the next batch.

Configuring full sync

See “Defining the CDP gateway ” on page 206. for information on how to configure full sync on NetBackup version 10.0 onwards.

In NetBackup 9.1, you can configure the number of concurrent full sync operations by specifying a value for the *CCT_MAX_FULL_SYNC_REQS* parameter, in the `nbcct.conf` file. For example, CCT_MAX_FULL_SYNC_REQS=7

Monitoring CDP jobs

More information is available on monitoring jobs in the web UI.

NetBackup dashboard

CDP follows the same job hierarchy as the traditional NetBackup agent for VMware. Protection starts with the job of discovering the VM and its attributes. A child job called Preparing for Backup follows it. This child job determines the changed blocks based on previous images and current data available on the gateway. A backup job, to move data from the CDP gateway to the destination storage unit, follows the child job.

If there is not enough space for each VM, on the gateway, the backup image may not be fully recoverable. Such images are referred to as partial non-recoverable images and are not available to restore from the web UI. But the subsequent backup jobs, create recoverable backup images. If an image is non-recoverable, NetBackup triggers a backup job automatically when it receives consistent data from the ESXi.

Viewing notifications

For most CDP activities, you can see notifications in the web UI. These notifications are helpful to know how the I/O tapping on the gateway performs. You can see notifications when things have stopped working or if any action is required from your side. The following are some important scenarios when you can see notifications:

- While backing up data. When a backup job moves data from the staging area to backup storage.
- VM full sync has started/suspended/resumed/done.
- Partial image is generated.
- No space left in the staging area for storage.
- When there is an error while writing in-memory data to the staging area.

Here are some notifications:

Table 14-2 Viewing notifications

Message	Scenario	Severity	Priority
Temporarily disconnecting from the IO filter to the Continuous data protection service on the gateway. Either the allocated staging area is almost full, or the memory usage is at maximum.	The staging space allocated to CDP is almost full, and CDP service temporarily disconnects from the I/O filter. This may also happen, if backup jobs cannot move data from the CDP gateway staging database to the backup storage. Check the backup job failure reasons and STU's underlying storage.	Critical	High
Input/Output error occurred for the VM: <uuid>	CDP service cannot perform I/O on staging location due to many reasons, like the underlying disk snapped out of storage, the file system going into read-only mode, and so on.	Error	High
Terminating the Continuous data protection service, as the staging area memory is full.	If the staging space is less than 1 GB, CDP raises this error and terminates the service.	Critical	High
Data storage quota is full for the VM: <uuid>, bearing jobid: ${jobid}. Moving data to the backup storage.	During the VM's data transfer, if the total data crosses the configured VM quota, then a backup job is triggered to move the staging data to the backup destination.	Info	Low
Cannot move data to the backup storage, for the VM: <uuid>. The storage quota for the VM is full.	Data movement from the gateway to the backup location failed.	Error	High
Full sync started for the VM: <uuid>.	Initiated the full sync process for this VM.	Info	Low
Full sync resumed for the VM: <uuid>.	Full sync for the VM is resumed after some unexpected interruption.	Info	Low

Table 14-2 Viewing notifications *(continued)*

Message	Scenario	Severity	Priority
Full sync completed for the VM: <uuid>.	The initial full sync for VM is complete.	Info	Low
Full sync suspended for the VM: <uuid>.	Full sync operation fails, for some reason like, a network glitch.	Info	Low
Backup image generated for the VM: <uuid> is not recoverable.	When a VM sync is in progress, if the VM quota is reached, a backup job is triggered. When the backup job is completed, the image may not be recoverable, as NetBackup is moving the intermediate data generated on the guest VM.	Info	Low

Viewing jobs

CDP uses the activity monitor to display the following job information:

- Parent backup job - discovery job to discover the VM information.
- Preparing for backup - identify the point-in-time data for the VM.
- Backup - move data from the staging path to the backup storage.

Using accelerators with CDP

CDP for VMware exclusively supports accelerator-based backup. So, CDP needs accelerator-compliant storage units based on MSDP or OST-based storage.

Force rescan

Force rescan enhances safety, and establishes a baseline for the next accelerator backup. This feature protects against any potential damage like failure of checksum verification on the data in the staging area.

When you use accelerator-based forced rescan, it clears the data on the CDP gateway staging area. So, any corrupted data is replaced with fresh data synced from the ESXi server. Note that the first backup job triggered by forced rescan may not have all data needed for a recoverable image. As data becomes available, the subsequent backups are triggered automatically, making the images recoverable.

Recommendations for using forced rescan:

- Do not trigger force rescan for the VMs which are turned off.

- If the staging location memory is full, you can see a notification in the UI. Initiate the force rescan only when sufficient memory is available at the staging location.

To manually trigger the backup with force rescan run the following command in the command prompt or the Linux terminal:

```
bpbackup -i -p policyname -s <schedulename>
```

NetBackup creates a schedule named `ForcedRescan` for every protected VM.

Recovering CDP protected VMs

VMs protected by NetBackup CDP for VMware have the same backup image format as the NetBackup agent for VMware. So, all recovery operations are the same as the NetBackup agent for VMware.

Here are some minor differences:

- Agentless single file recovery is supported only if MSDP is configured for instant access.
- Recovery from the vCenter plug-in is not supported.
- Cannot restore VMs from CDP-based backup images through Java UI.

Web UI does not allow recovery of the images shown as partial and non-recoverable. You can restore them using NetBackup API. However, the VMs may not start after the recovery.

Some limitations of CDP

Here are some limitations of CDP:

- NetBackup features like Intelligent policy and Backup now, and Roll back instantly from web UI are not supported.
- CDP for VMware and Veritas Resiliency Platform do not work together for the same VM. However, both products can protect different VMs on the same vCenter cluster.
- CDP does not support any standalone ESX, which is not managed by any VC. An ESXi which is not part of any ESXi cluster but is managed by VC, is also not supported.
- You must turn on the VMs before subscribing them to a CDP-based protection plan, and also for the first full backup.

- After subscribing a VM for CDP backup policy, if any disk from the VM is removed or a new disk is added, the subsequent backups fail. In such cases, unsubscribe the VM from CDP protection, and subscribe it again.
- Due to VMware limitation, if you try to protect a VM using the NetBackup agent for VMware and CDP, both at the same time, the backup operation fails with an error or the operation might crash with symbols from VDDK.

Troubleshooting for CDP

VAIO stops sending data to the CDP gateway

This happens when the IOFilter encounters problem and hence enters into NOOP (Non-Operational) mode.

Possible reasons:

- IOfilter encountered a problem with the datastore.
- IOfilter encountered a problem while reading from vmdk on the ESXi server.

Workaround:

Remove the VTSTAP policy from all the disks of protected VMs and reattach.

Error: Storage policy is not detached from one or more virtual disks of the virtual machine.

This happens when the storage policy is not detached from all the virtual disks of the VM. The next backups fail with error code 156.

Workaround:

Remove the Veritas I/O filter-based storage (vtstap) policy from all the disks of the VM that CDP protected previously. You can do this operation on the vCenter.

Error: Failed to retrieve or parse the version of Veritas IO filter.

You may get this error when trying to subscribe one or more VMs to the CDP protection plan. Occurs when the CIM server service on the ESXi server is non-responsive.

Workaround:

Restart the CIM server service on the ESXi server and retry the VM subscription to the CDP protection plan. You can find the CIM server service of the ESXi server, under Configure > Services section of the ESXi.

nbcctd service goes to an inconsistent state. Cannot configure the CDP gateway.

Possible reasons:

- When you mount a read-only file system and provide its path in the CDP gateway configuration, the service is configured, but the gateway fails to start.
- When you try to configure the gateway again, by giving a read/write path, the service still fails to start.

Workaround: Retry the operation after you remove the `nbcct` directory from:

- `<NBU installation path>/netbackup/nbcct` in NetBackup 9.1.
- `<staginglocation>/nbcct` in NetBackup 10.0 and later.

CDP-based protection plan fails with the error: Storage policy is not attached to one or more virtual disks of virtual machine to be registered for IO tapping.

Possible reasons:

Currently, NetBackup supports only the vtstap policy as a storage policy for CDP. If you try to subscribe a VM using a hybrid storage policy (encryption + replication) it shows the error.

Workaround: Avoid using a hybrid storage policy (encryption + replication) for CDP-protected VMs.

CDP service does not start after the media server restart or mount path-related changes.

Possible reasons:

The configured staging area is unmounted post reboot or has an unsupported file system. For example, if you configure the CDP gateway using a supported mount like `/mnt/stage_area` and do not configure auto-mount. After a system restart, this path points to root file system, which CDP does not support, hence the CDP service (nbcctd) cannot start.

Workaround: Ensure that the staging area or the relevant disk mounts are remounted properly, whenever there are changes in the system related to unmount or system reboot.

VM gets unsubscribed in a powered off state and has I/O tapping policies attached to the VMDK. It gives a warning to remove storage policies and then unsubscribe.

Possible reasons:

While removing CDP protection, if the protected VM is powered off, the CDP gateway cannot get the required information of storage policies from VAIO. Hence, though the CDP protection is removed from the VM, the I/O tapping policies are still attached to the VMDK of that VM, it continues to tap the I/O and affect performance.

Workaround: Always detach the storage policy of the VMs before unsubscribing the VMs, irrespective of their powered on or off state.

Subscription to NetBackup protection plan fails, but the backup jobs keep dumping data in the staging area.

Explanation

Occurs when you protect a NetBackup primary server, using the same primary server's protection plan.

Workaround: We do not recommend protecting a NetBackup primary server using a protection plan made using the same primary server. If this error occurs, detach the Veritas storage policy from the NetBackup primary server VM, and unsubscribe the VM from the protection plan.

Cannot delete a CDP protection plan when the CDP gateway is unreachable.

Explanation:

CDP policy is not deleted after removing the entries in case of an unreachable host.

Workaround: The CDP protection plan subscription does not get removed, as we are not deleting the CDP policy before cleaning up the CDP host. So, we need to call the Delete policy API manually after calling the Delete CDP gateway API, to delete the entries of the unreachable gateway.

You can clean up an unreachable CDP gateway using the following API:

```
To DELETE CDP Gateway

URL : https://netbackup/config/cdp-gateway/force

HTTP Method : DELETE

Headers:

    Authorisation: Bearer <Token>

    Content-Type: application/vnd.netbackup+json;version=9.0;charset=UTF-8

To Delete Policy
```

```
URL : https://netbackup/config/policies/policy_name

HTTP Method : DELETE

Headers:

    Authorisation: Bearer <Token>
```

After the successful execution of the above two APIs, the mapping for the policy and the VM is still visible in the web UI. If you try to remove the protection of that VM through web UI, you can see an error message saying: **Subscription ID not found**. This is expected behavior.

CDP gateway update operation fails to restart the CDP service (nbcctd) on the gateway

Explanation: The CDP gateway update operation tries to restart the service. If stopping the service takes longer time than usual, then the update operation shows an error, indicating that the CDP service (nbcctd) failed to restart.

Workaround: In this case, check if the `nbcctd` service is running on the gateway. If the service is running, wait for it to shut down. To manually stop the service, use the command: `/usr/openv/netbackup/bin/nbcctd -terminate`. When the service has stopped, start it using the command `/usr/openv/netbackup/bin/nbcctd -X`.

Failed to get version from the Storage Platform Web Service(SPWS). Ensure that Nginx is running and configured correctly on the selected MSDP storage server.

Explanation: While creating a CDP protection plan to use universal share, if you select a storage device that does not have universal share capability, you get this error.

Workaround: You must select a storage device that has universal share capability.

Unsupported CDP gateway version with universal share. Minimum supported version is 10.2.

Explanation: While creating a CDP protection plan to use universal share, if you select a CDP gateway server lower than NetBackup version 10.2, you get this error.

Workaround: To use universal share, the CDP gateway version must be NetBackup version 10.2 or higher.

Chapter 15

Backing up virtual machines

This chapter includes the following topics:

Manually back up virtual machines

You can start a backup manually from a policy.

To manually back up virtual machines

1. Open the NetBackup web UI.
2. On the left, click **Protection > Policies**.
3. Select the policy and then select **Manual backup**.
4. Select the type of schedule for the backup.
5. Select the clients (virtual machines) to back up.

   If the policy was configured for automatic selection of virtual machines, the Clients list shows the VMware backup host rather than the virtual machines.

6 Click **Backup** to start the backup.

7 To see the job progress, click **Activity monitor**.

Note that your VMware backup request launches more than one job. The first job automatically creates and deletes snapshots. This job has a dash (-) in the Schedule column. The second job backs up the virtual machine files from the snapshot.

Trial backup for VMware

Trial backup for VMware lets you perform approximate validations for the backup configuration. Trial backup support for VMware is only available through an API call.

A trial backup operation for a virtual machine performs certain operations as:

- Discovery
- Validation of credentials
- Snapshot
- Mapping of a few files (if you have selected the enable file recovery option in the policy).
- Test the transfer of limited bytes (to the media server) of data per disk of a virtual machine.

You can use the following API with the input parameter `trialBackup: true` to initiate a trial backup for a VMware policy. Before you use the trial backup API, the required policy configuration must be defined in the policy created.

API : /admin/manual-backup

Method : POST

Example: **Request body**

```
{
 "data": {
                "type":"backupRequest",
                "attributes": {
                    "policyName": "vmware_test",
                    "trialBackup" : true
                }
     }
}
```

The backup job displays as **Trial Backup**.

Note the following:

- The parameter `trialBackup` is optional and its default value is false.
- A trial backup is not recoverable. The **Image Clean up** job does the cleanup of test data which is transferred during the trial backup operation.
- The trial backup functionality is only supported for a VMware policy of VADP (VMware vStorage API for Data Protection) based backups. It does not validate application integration with Microsoft Exchange, SharePoint, and SQL server.
- VMware agent resource limits apply to trial backup.
- A trial backup consumes resources like querying the vCenter for discovery and there can be short-lived snapshots.
- A trial backup doesn't affect existing schedules or incremental backup chains. However, it may affect VMware resources like the vCenter for discovery and snapshot.

Using the Activity monitor to monitor virtual machine backups

You can use the NetBackup Activity monitor to keep track of the VMware virtual machines that a policy backs up.

To monitor virtual machine backups

1. On the left, click **Activity monitor**.
2. Note each job as it appears by row.

   If the policy selects virtual machines automatically (based on a query), the backup consists of three generations of jobs:

   - The first job discovers the virtual machines. This job is labeled **Backup**. (This job is unique to policies that use a query to select virtual machines.)
   - The discovery (**Backup**) job starts a child job to take a VMware snapshot of the virtual machine. A snapshot job is started for each virtual machine. Each of these jobs is labeled **Snapshot**.
   - Each snapshot job starts a child job to back up the virtual machine. A backup job is started for each virtual machine. Each of these jobs is labeled **Backup**. The job flow is as follows:

     ```
     discovery job --> snapshot job --> backup job
     ```

3 To trace the discovery job to the virtual machine backup jobs, note the **Job ID** and the **Parent job ID** columns.

4 Click on the job and click **Details** tab.

See "Limit jobs per policy on the Attributes tab (for VMware)" on page 90.

Restarting jobs individually in the Activity monitor

If the policy automatically selects virtual machines for backup, you can restart the virtual machine jobs individually. This feature is handy if the policy backs up a large number of virtual machines: you can restart one or more of the jobs individually rather than re-running the entire policy.

To restart jobs individually in the Activity monitor

- In the Activity monitor, locate the job and click **Actions > Restart**.

  In some cases, to restart a child job you may have to restart its parent job.

Viewing NetBackup activity in vSphere Client (HTML5)

In VMware vCenter, NetBackup can record the backup activity for virtual machines. You can view the events in vSphere Client (HTML5) at the level of any parent object (such as folder, datacenter, cluster, or host). You can also view the events for the virtual machine.

Note: The NetBackup plug-in is not required.

Make sure that the policy's **Post vCenter events** option is enabled:

See "VMware - Advanced attributes" on page 96.

To view backup events and the last backup time in vSphere Client

1 Open the vSphere Client (HTML5).

2 Go to **Home > Events**.

3 In the **Event Console**, enter `NetBackup` in the search field.

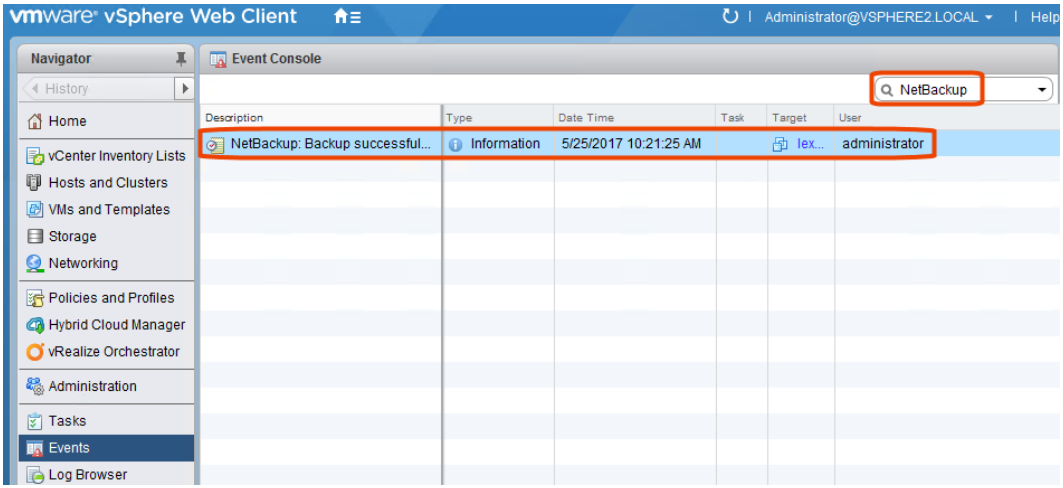

Each NetBackup event includes the following details:

Description	Shows the NetBackup operation that succeeded or failed, and includes policy details and duration. When you click on the row of the event, the **Event Details** pane shows the same information.
Type	The types are **Information** for a successful operation, and **Error** for a failed operation.
Date Time	The date and time of the event.
Task	Not used.
Target	The virtual machine that was backed up. Click on the virtual machine link to see the **Summary** tab. **Note:** On the **Summary** tab in vSphere 6.5, under **Custom Attributes**, the timestamp of the virtual machine's last backup is listed on the `NB_LAST_BACKUP` attribute.
User	The user that ran the backup.

4 To view backup events for a particular object (such as ESX host or VM), select the object and do the following:

- Click the **Monitor** tab.

- Click **Task & Events**.
- Click the **Events** view.
- Enter `NetBackup` in the search field.

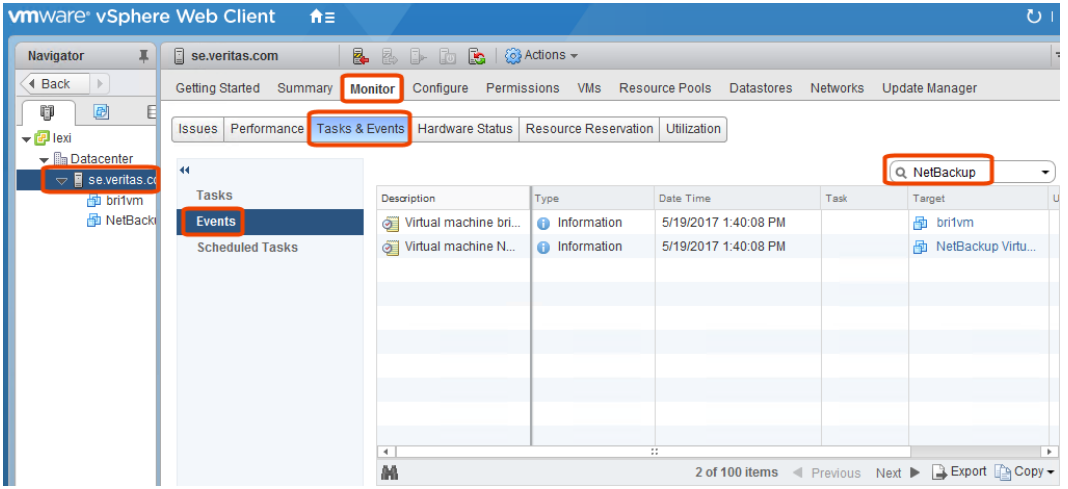

Chapter 16

VM recovery

This chapter includes the following topics:

Restore notes and restrictions

Before you begin the restore, note the following:

- Cross-platform restore of individual files is not supported. You can restore Windows files to Windows guest operating systems but not to Linux. You can restore Linux files to supported Linux guest operating systems but not to Windows. In other words, the restore host must be the same platform as the files that you want to restore.
  See “About restoring individual VMware files and folders” on page 249.
- If you have back-level hosts in your environment, note the following about mixed-level backups and restores: The recovery host must be at the same or a later NetBackup release level as the backup host. For example, you cannot use a NetBackup 8.x recovery host to restore a virtual machine that was backed up by a NetBackup 9.x backup host.
- Unless a NetBackup client is installed on the virtual machine, you must do the restore from the NetBackup primary server. Or, perform a VMware agentless restore or create an instant access VM for the restore.
- To restore files to the original virtual machine location, the destination must be specified as the virtual machine's host name (not display name or UUID).

- To restore directly to an ESX server, the name that is specified for the restore must match the ESX server's official host name. The name must be in the same format in which it is registered in DNS and in the VMware server (whether short or fully-qualified).
  See "Add VMware servers" on page 66.
- If the VM's display name was changed after the VM was backed up, the pre-recovery check may fail when you click **Start Recovery**:

  ```
  VM exists overwrite -Failed. Vmxdir for VM exists
  ```

  You can ignore the error and click **Start Recovery**, but note: The restore may succeed but the folder that contains the vmx file for the newly restored VM has a different name than the vmx folder of the existing VM. VMware does not rename this folder when the VM is renamed, but continues to use the existing folder.

  As an alternative, restore the VM to a different location.
- A virtual machine template cannot be restored to a standalone ESX server. Because templates are a feature of vCenter servers, you must restore the template through vCenter. If you restore a template to a standalone ESX server, the template is converted to a normal virtual machine and is no longer a template.
- NetBackup supports backup and recovery of VMware NVRAM files and the vTPM devices that are associated with virtual machines.
  - A NetBackup 8.3 or later backup or recovery host is required for NVRAM and vTPM protection. Supported recovery methods include Full VM recovery and VMware Instant Recovery.
  - NetBackup does not support the backup or restore of NVRAM and vTPM for the virtual machines whose display names begin with a period ('.'). An existing VMware limitation prevents downloading or uploading data store files beginning with a period ('.') to a virtual machine's working directory as these appear as hidden files.
- Restore of individual files from a backup of the full virtual machine is not supported if the virtual machine contains Storage Foundation Volume Manager volumes.
- To restore Windows NTFS-encrypted files individually, you must install a NetBackup client on the virtual machine.
  See "NetBackup for VMware best practices" on page 306.
- VMware does not support the restore of virtual machines directly to an ESX 5.x server that vCenter manages. To restore the virtual machine, select the vCenter server as the destination.

As an alternative, you can set up an independent ESX server to be used for restores. You must add NetBackup restore credentials for that ESX server by means of the **VMware restore ESX server** server type.
See "Add VMware servers" on page 66.

- The APIs in VMware's Virtual Disk Development Kit (VDDK) contain the following limitation: The maximum write speed during virtual machine restore is roughly one third of the hardware's maximum speed.
- If a virtual machine had vmdk files in different directories in the same datastore, note: When the virtual machine is restored to the original location its vmdk files are restored to a single directory, not to the original directories. (This behavior follows current VMware design.)
  As a workaround, do the following: Remove the vmdk files from the restored virtual machine, move the files to their respective directories, then re-attach the moved files to the virtual machine.
- If the original VM contains encrypted vmdk files, after restoring the full VMware virtual machine to the original location or after an in-place disk restore, the restored disks may not be compliant to the VM encryption policy. Therefore, the restored VM must be reconfigured manually to comply with the policy. Otherwise, the virtual disks of the restored VM might be left in an unencrypted state.
- When restoring large files, make sure that no snapshots are active on the destination virtual machine. Otherwise, the files are restored to the VMware configuration datastore, which may be too small to contain the files you want to restore. In that case, the restore fails.
  The configuration datastore (sometimes called the vmx directory) contains the configuration files that describe the virtual machine, such as `*.vmx` files. Note that active snapshots of vmdk files are also stored on the configuration datastore.
- If you cancel the virtual machine restore before it completes, the not-fully-restored virtual machine remains at the target location. NetBackup does not delete the incomplete virtual machine when the restore job is canceled. You must manually remove the incomplete virtual machine.
- If the virtual machine display name contains unsupported characters, the backup may succeed but the restore fail. To restore the virtual machine, you must change the display name to contain supported characters only and retry the restore.
  See "NetBackup character restrictions for the Primary VM identifier" on page 42.
- NetBackup for VMware does not support individual file restore by means of client-direct restore.
- On a restore, NetBackup recreates the linking between a hard link and its original file only in this case: The link file and its target file are restored in the same job.

If each file is restored individually in separate restore jobs, they are restored as separate files and the link is not re-established.

- If you restore a VM in vCloud to an expired vApp, the vApp is automatically renewed and added back into the vCloud organization. If the expired vApp contained other VMs, all those VMs are also removed from the expired list and added to the organization.
  Note that in vCloud Director, an expired vApp must be renewed before you can import a VM into that vApp.
- With a remote connection from a Windows Java GUI that uses the English locale, the restore of files that have non-ASCII characters may fail.
  See the following tech note for further information on how to restore the files:
  https://www.veritas.com/docs/100022268
- In VMware for Replication Director and Integrated Snapshot Management policies, if you configure SLP as combination of Snapshot and Index from Snapshot (IFS), then the restore of files on XFS formatted volumes and partitions is not supported via NetBackup Java UI, use NetBackup web UI.
- In VMware vSphere 6.0 U1b and later, a full restore of a virtual machine may trigger an alarm if the original VM was not deleted. The alarm is a VM MAC address conflict alarm. This VMware alarm behavior is by design. If there is a MAC address conflict, VMware eventually changes the MAC address of the new VM. If you do not want to receive alarms, disable the VM MAC address conflict alarms in vCenter.
- See “NetBackup for VMware: notes and restrictions” on page 35.

Restore notes and restrictions on Linux

This topic relates to the restore of files from a NetBackup backup of a VMware virtual machine that runs Linux.

Before you begin the restore, note the following:

- Cross-platform restore of individual files is not supported. You can restore Linux files to supported Linux guest operating systems but not to Windows.
- To migrate an ext2 or ext3 file system to ext4, note: Make sure to follow the instructions under "Converting an ext3 file system to ext4" on the following page of the Ext4 wiki:
  https://ext4.wiki.kernel.org/index.php/UpgradeToExt4
  If you do not follow these instructions, data in a newly created ext4 file is not promptly flushed from memory to disk. As a result, NetBackup cannot back up the data of recently created files in the ext4 file system. (The NetBackup snapshot captures the file as zero length.)

As a workaround for the file systems that were not correctly migrated, do one of the following:

- Run the Linux sync command on the ext4 file system before starting each backup.
- Make sure that snapshot quiesce is enabled in the Linux guest OS. Contact your operating system vendor and VMware for additional information.

- For Linux virtual machines, NetBackup cannot restore individual files from software RAID volumes. The files are restored when you restore the entire virtual machine.
- The Linux ext4 file system includes a persistent pre-allocation feature, to guarantee disk space for files without padding the allocated space with zeros. When NetBackup restores a pre-allocated file (to any supported ext file system), the file loses its preallocation and is restored as a sparse file. The restored sparse file is only as large as the last byte that was written to the original file. Note also that subsequent writes to the sparse file may be non-contiguous.
- NetBackup supports backup and restore of Linux LVM2 volumes, including individual file restore from an LVM2 volume. Note however that NetBackup does not support individual file restore from a snapshot that was created by means of the snapshot feature in LVM2. If an LVM2 snapshot exists at the time of the backup, the data in the snapshot is captured in the backup. The data can be restored along with the rest of the virtual machine data when you recover the entire virtual machine.
- NetBackup supports backup of Linux FIFO files and socket files. NetBackup does not support restoring FIFO files and socket files individually. FIFO files and socket files can be restored along with the rest of the virtual machine data when you recover the entire virtual machine.
- When you restore Linux files individually to an NFS-shared device on a Linux virtual machine, NetBackup can only restore the file data and attributes. The extended attributes cannot be restored to NFS-shared devices.
- For a virtual machine that is running a Linux guest operating system: When you restore a virtual machine, the ESX server may assign the virtual machine a new (virtual) MAC address. After you restart the virtual machine, you may have to configure its MAC address. For instance, the original MAC address of the virtual machine may be in a configuration file that has to be updated.
  Refer to your VMware documentation for more details.
- For Linux, additional notes apply.
  See “NetBackup for VMware: notes on Linux virtual machines” on page 39.

Recover a full VMware virtual machine

You can recover a VM to its original location where it existed when it was backed up or to different location. You can choose to recover from the default copy of the backup image or from an alternate copy, if one exists. The default copy is also known as the primary copy.

To recover a VM

1. On the left, click **Workloads > VMware**.
2. Locate and click on the VM.
3. Click the **Recovery points** tab. In the calendar view on the left, select the date on which the backup occurred.

   The available images are listed in rows with the backup timestamp for each image.
4. (Conditional) Select the **Allow the selection of recovery of points that are malware-affected** to be able to recover malware affected images.

   Note: This option is only available for the recovery points which contains malware-affected images. And, it is enabled only for users with that have an RBAC role with the necessary permissions.
5. For the image that you want to recover, select one of the following image recovery options:
   - **Recover**
     Recover from the default copy of the backup image. This option is displayed if only one copy exists.
   - **Recover from default copy**
     Recover from the default copy of the backup image. This option is displayed if more than one copy exists.
   - ***nn* copies**
     Recover from the default copy or a different copy of the backup image. NetBackup allows up to ten copies of the same backup image. All available copies are displayed when you select this option. For each copy, the **Storage** name, **Storage server**, and the **Storage server type** are displayed.
6. Click **Restore virtual machine** for the copy that you want to recover.
7. On the **Restore to** tab, do the following:
   - Review the **Restore to** values.
     The default values are populated from the backup image of the VM.

- To recover to an alternate location, change the values for **ESXi server or cluster**, **Folder**, or **Resource pool or vApp**.
- Select the appropriate option for **Use datastore or storage policy**.
- Click **Next**.

8 The restore options that display next depend on if you chose to restore to the selected datastore or to use a storage policy.

See “Recovery options” on page 261.

See “Storage policy” on page 233.

(Restore to selected datastore) Review or change the **Advanced** options:

See “Advanced recovery options” on page 233.

See “Advanced recovery options: Format of restored virtual disks” on page 234.

See “Advanced recovery options: Transport mode” on page 235.

9 Click **Next**.

10 NetBackup performs pre-recovery checks that include verifying the credentials, appropriate paths, connectivity, and determining if the datastore or datastore cluster has available space.

11 Resolve any errors.

You can choose to ignore the errors. However, the recovery may fail.

12 Click **Start recovery**.

Click the **Restore Activity** tab to monitor a job's progress. Select a specific job to view its details.

Recovery options

Allow overwrite of existing virtual machine	NetBackup deletes any VM with the same display name that exists at the destination, before the recovery starts. Note that, NetBackup deletes any VM with the same display name, it may not be the same VM, but another VM having the same display name. For VMware Cloud Director recovery, this option is not displayed if you choose to restore to **vSphere** (not **VMware Cloud Director**).
Power on after recovery	Automatically turns on the VM when the recovery is complete.
Recovery host	Indicate the host that you want to use to perform the recovery. By default, the recovery host is the one that performed the backup.

Storage policy

The storage policy settings are available for a virtual machine restore when you select option **Use a storage policy to select datastore** on the **Recovery target** page.

Virtual machine storage policies control which type of storage is provided for the virtual machine. You can apply one storage policy to the entire VM or you can apply different storage policies to the VM home directory or virtual disks.

Apply to whole virtual machine

Select storage policy	Select a storage policy to apply for the whole virtual machine from the list of all storage policies that are associated with the selected vCenter server.
Datastore or datastore cluster	Select a datastore that is compatible with the selected storage policy.

Customize virtual machine

Virtual disk	Lists the VMs home directory and or virtual disks that were captured at backup time and associated storage policy information.
Storage policy	Select a storage policy to apply to the VM home directory or the virtual disk from a list of all storage policies that are associated with the selected vCenter server.
Datastore or cluster or path	Select a datastore that is compatible with the selected storage policy.

Advanced recovery options

Create a new BIOS UUID	Restores the VM with a new BIOS UUID instead of the original BIOS UUID.
Create a new instance UUID	Restores the VM with a new instance UUID instead of the original instance UUID.

Remove backing information for devices

For example, this option restores the VM without restoring any ISO file that was mounted when the VM was backed up.

If this option is disabled, the recovery might fail if the backing information is not longer available for devices, such as DVD/CD-ROM drives, or serial or parallel ports.

Remove original network configuration

Removes the NIC cards from the VM. Note that for network access, the restored VM requires network configuration.

Enable this option if:

- The network connections on the destination virtual machine have changed since the backup was made.
- The original virtual machine still exists and a duplicate VM may cause conflicts.

Remove tag associations

When this option is selected, NetBackup does not attempt to restore tag associations when it restores the virtual machine. If the option is disabled, NetBackup attempts to restore all tag associations from the backup. If NetBackup cannot restore one or more of the tag associations, the restore exits with a NetBackup status code `1`.

See "Notes and limitations for the backup and restore of VMware tag associations" on page 47.

Retain original hardware version

Restores the VM with its original hardware version (such as 4). It retains the original version even if the target ESXi server by default uses a different hardware version (such as 7 or 8). If the target ESXi server does not support the virtual machine's hardware version, the restore may fail.

If this option is disabled, the restored virtual machine is converted to the default hardware version that the ESXi server uses.

Advanced recovery options: Format of restored virtual disks

Original provisioning

Restores the VM's virtual disks with their original provisioning.

Thick provisioning lazy zeroed

Configures the restored virtual disks in the thick format. The virtual disk space is allocated when the disk is created. This option restores the populated blocks, but initializes vacant blocks with zeros later, on demand.

Note: If the vmdk is completely written, VMware automatically converts a lazy-zeroed disk to **Thick provisioning eager zeroed**.

Thick provisioning eager zeroed	Configures the restored virtual disks in the thick format. Restores the populated blocks and immediately initializes vacant blocks with zeros (eager zeroed). Creation of the virtual disks may take more time with this option. However, if the restore occurs over a SAN, the eager zeroed feature may speed up the restore by reducing network communication with the vCenter server.
Thin provisioning	Configures the restored virtual disks in the thin format. Restores the populated blocks but does not initialize vacant blocks or commit them. Thin provisioning saves disk space through dynamic growth of the vmdk file. The vmdk files are no larger than the space that the data on the virtual machine requires. The virtual disks automatically increase in size as needed. **Note:** If the vmdk is completely written, VMware automatically converts a thin disk to **Thick provisioning eager zeroed**.

Advanced recovery options: Transport mode

The **Transport mode** specifies the mode to use for backups or how to read the data from the datastore. For more information on transport modes, see the vendor documentation for your virtualization environment.

Note the following when you select a transport mode:

- The SAN mode is not supported for the virtual machines that use VMware Virtual Volumes (VVols).
- For the hotadd mode, the virtual machines that use VVols and the backup host (hotadd) virtual machine must reside on same VVol datastore.
  See "Notes on the hotadd transport mode" on page 45.

Restoring VMware virtual machine disks

Use this procedure to use the NetBackup web UI to restore VMwrae virtual machine disks.

See "About VMware virtual machine disk restore" on page 236.

To restore virtual machine disks

1. Open the NetBackup web UI.

   Note that to perform this procedure you must have the RBAC Administrator role or a role with similar permissions.
2. On the left, click **Recovery**.
3. On the **Regular** recovery card, click **Start recovery**.

4. For the Policy type, select **VMware**.
5. From the **Restore type** list, select **Virtual disk restore**.
6. For the **Source client** select the VMware virtual machine that was backed up.

   Click **Select client** to search or browse for a virtual machine. Use this option to locate a virtual machine in a large, multi-layered virtual environment.

   Or enter the type of name that was selected for the **Primary VM identifier** option on the policy **VMware** tab. For example, if the**Primary VM identifier** option is set to VM host name, enter the virtual machine's host name.

   If any backups for the client are within the specified date range, NetBackup populates the right pane with the information about the most recent backup.
7. To restore disks from a backup other than the most recent, follow these steps:
   - Next to **Date range** click **Edit**. Then select **Use backup history**.
   - Select the wanted backup and click **Apply**.
8. Select the client. Then click **Next**.
9. On the right, click **Virtual disk** or **File system**. Select the virtual disk or the file system.

   See "Selecting virtual disks or file systems" on page 237.

   Click **Next**.
10. Select the **Recovery options**.

    See "Recovery options for virtual machine disks" on page 238.
11. Select the **Storage target**.

    See "Storage target restore options" on page 239.
12. Review the recovery options and details. Then click **Start recovery**.

About VMware virtual machine disk restore

The following are the general support requirements for virtual disk restore.

- Sufficient storage must exist for the restore.
- NetBackup does not support the following virtual machine disk restores:
  - From NetBackup Replication Director for VMware backups.
  - To templates. However, virtual disks from a backup of a VM template can be restored to a virtual machine.

NetBackup supports the restore of individual VMware virtual machine disks to the following destinations:

To the original VM	You can restore the disks to the same VM from which the disks were backed up. You can either overwrite the original disks or attach the virtual disks without overwriting the original disks. NetBackup creates a temporary VM to which it restores the virtual disks. Then, NetBackup attaches the virtual disks to the existing, target VM. Finally, NetBackup deletes the temporary VM after the disk or disks are attached successfully. A special case, called in-place disk restore, replaces all disks of an existing VM with the data in its backup. Raw devices (RDMs) and independent disks are not replaced or deleted. For In-place Disk Restore, the disks are restored to the same disk controller configuration acquired at the time of backup.
To a different VM	You can restore the disks to a different VM. NetBackup creates a temporary VM to which it restores the virtual disks. Then, NetBackup attaches the virtual disks to the existing, target VM. Finally, NetBackup deletes the temporary VM after the disk or disks are attached successfully. You can also perform an in-place disk restore to a different VM.
To a new VM	NetBackup creates a new virtual machine and restores the specified disks to the new VM. The new VM is intended to be a container for the restored disks. It does not have enough resources to run most operating systems. After the restore, you should attach the restored virtual disks to a VM that can support them and then delete the restore VM.

Selecting virtual disks or file systems

The **Virtual disks** page shows all of the virtual disks that were in the VM at back-up time, even those that were excluded from the backup. By default, NetBackup displays the virtual disks. To select file systems, click **File system**.

- **Virtual disks**
  Select the disks that you want to restore.
  If you select all disks, NetBackup selects only those disks that were included in the backup.
- **File system**
  Select the wanted file systems. When you select a file system, NetBackup selects the virtual disks on which the file system resides.
  If you select all file systems, NetBackup selects only those file systems that were included in the backup.

Recovery options for virtual machine disks

You can configure the following recovery options when you restore a virtual machine disk.

Table 16-1 Restore options for virtual machine disks

Option	Description
Restore to	The destination VM for the restore, as follows: ■ **Original virtual machine**. NetBackup restores the selected disks to the VM from which they were backed up. You can select whether to overwrite the existing virtual disks in the **Storage target** settings. See "Storage target restore options" on page 239. ■ **Alternate virtual machine**. NetBackup restores the selected disks to a different VM than the original. Select the destination VM for the restored virtual disks. **Last updated** indicates the date and time at which the **Virtual machine server** provided the information to NetBackup. Click **Discover** to update the virtual machine server details. ■ **New (temporary) virtual machine**. NetBackup creates a new virtual machine and restores the selected disks to the new VM. The new VM is intended to be a container for the restored disks. It does not have enough resources to run most operating systems. After the restore, you should attach the restored virtual disks to a VM that can support them and then delete the temporary VM.
Power on after recovery	Select this option to have the recovered virtual machine automatically turned on when the recovery is complete.
Recovery host	The host that performs the restore. If not specified, NetBackup uses the backup host value from the backup image.
Media server	You can use this option to select a media server that has access to the storage unit that contains the backup image. An example of such an environment is a Media Server Deduplication Pool (MSDP) with multiple media servers. Note: If the storage unit that contains the backup image is not shared with multiple media servers, this option is grayed out.

Table 16-2 Advanced restore options

Transport mode	The transport modes to use for the restore. By default, NetBackup selects the transport mode that was used for the backup. Alternatively, you can select the wanted transport modes and arrange them in the priority order that you want. If all methods fail, the restore fails.

Table 16-2 Advanced restore options *(continued)*

Delete restored staging VM on error	Whether to delete the temporary VM if the disk attach operation fails. If you disable this option and the disks are not successfully attached to the target VM, you can access the data on the temporary VM. Applies only to **Original virtual machine** or **Alternate virtual machine**.
Wait time for VM shutdown	The restore process shuts down the target virtual machine before it attaches the disk or disks. The duration of the shutdown operation depends on the VMware workload. Use this parameter to specify how long the restore process should wait for shutdown before giving up on restore. Applies only to **Original virtual machine** or **Alternate virtual machine**.

Storage target restore options

The **Storage target** page includes additional restore options for the virtual disks to restore.

Table 16-3 Storage target restore options

Option	Description
Apply to whole virtual machine	The following settings apply to the whole virtual machine. ■ **Overwrite all virtual disks** Whether to overwrite the existing virtual disk or disks on the target VM. If this option is enabled, overwrite the original virtual disk and retain the disk UUID. If the option is disabled, restore the virtual disk to the target VM as a new disk; VMware assigns a new UUID to the disk. ■ **Datastore** The name of the **Datastore** or the datastore cluster that is the destination for the restore. Click **Search** to select a different datastore or datastore cluster. ■ **Restored virtual disks provisioning** The default disk provisioning for all of the disks to restore: **Thin provisioning**, **Thick provisioning lazy zeroed**, or **Thick provisioning eager zeroed**.
Customize virtual machine	Select **Customize virtual machine** to override the global settings and change the values for individual virtual disks.

Chapter 17

VMware agentless restore

This chapter includes the following topics:

About VMware agentless restore

The agentless restore lets you restore individual files and folders to virtual machines where the NetBackup client is not installed. By using VxUpdate, NetBackup can deploy the recovery tool to the virtual machines, restore files and folders, and perform the required cleanup. NetBackup does not require a connection to the target virtual machine to recover the files. All recovery is handled through the ESX server using VMware vSphere Management APIs.

A video is available that describes NetBackup VMware agentless restore:

VMware agentless recovery video

Overview of the agentless restore process

1. The NetBackup primary server receives input from either the NetBackup web UI or the Agentless Recovery API. The input is the files and folders for restore along with the credentials for the target virtual machine. These credentials must have administrator, root, or sudo privileges.
2. The primary server sends the requested data to the restore host.

3. The restore host confirms that it has the necessary VxUpdate recovery package to perform restore. If it's not available, the restore host downloads the required package from the primary server using VxUpdate.
4. The restore host pushes recovery tool to virtual machine using the vSphere management API.
5. The data stream containing the user-selected files and folders is staged in a vmdk that is associated with a temporary virtual machine. Cohesity creates the temporary virtual machine for the agentless restore.
6. The vmdk that NetBackup created on the temporary virtual machine is attached to the target virtual machine.
7. The recovery tool is invoked and the files and folders are recovered.
8. NetBackup performs the necessary cleanup. All temporary files and objects that are created as part of the process are deleted or removed. Among the objects that are deleted and removed are the recovery tool, the temporary virtual machine, and the staging vmdk.
9. The job is finished.

Prerequisites and limitations of VMware agentless restores

Prerequisites

The following prerequisites exist for VMware agentless restores:

- You must provision VxUpdate packages for all platforms for which you have virtual machines where you want to perform agentless recovery.
- You must have credentials with administrator, root, or sudo permissions for the target virtual machine.
- The target VM is where the files are recovered. It must be powered on and have the latest version of VMware Tools installed.
- The target VM should have at least one Paravirtual Controller with available LUNs. Or, available space for a Paravirtual SCSI Controller.
- To use non-root credentials on a Linux target VM it must have sudo installed and the `/etc/sudoers` file configured so that the user has the following permissions:

```
username ALL=(ALL) NOPASSWD: /bin/tar, SETENV:
/usr/openv/tmp/rt/netbackup/bin/nbtar_rt
```

or

```
username ALL=(ALL) NOPASSWD: ALL
```

- The default staging location on the target VM is `%TEMP%` or `%TMP%` for Windows and the tmp directory (`/tmp`) for Linux.
- The staging location must exist on the target VM file system.
- If you want to allow the use of instant access for recovery of the files and folders, the recovery point must support instant access.
  See "Create an instant access VM" on page 185.

Limitations

The following limitations exist for VMware agentless restores:

- Agentless restores to Windows target VMs can fail if you use an account other than the built-in **Administrator for Windows Guest OS** account as the **Target VM Credentials**. The restore fails because **Run all administrators in Admin Approval Mode** is enabled. More information is available:
  https://www.veritas.com/content/support/en_US/article.100046138.html
- VMware agentless restores can only be used for the restore of files and folders.
- In some instances, when you perform an agentless restores, orphaned VMs starting with `NB_` are left behind. Using the ESX server credentials to perform the restore on the target VM even though the vCenter manages the ESX server can cause this condition. This condition is a known limitation of VMware. To resolve the problem, register the vCenter in NetBackup and use vCenter credentials for backups and restores. The orphaned VMs starting with `NB_` can be removed from inventory manually by logging into the vCenter using VMware vSphere Client.
- Restore job fails if NetBackup is unable to use the staging directory. This directory is specified in the `TMP` or `TEMP` environment variable.
- Restore job fails if NetBackup does not have sufficient privileges to the staging directory. Or, if there is insufficient space in the staging directory.
- If you select **Flatten existing directory structure** and **Overwrite existing files** options, you risk an incorrect restore if it contains multiple files with the same file name. In this case, the last file that is restored is the one that is present when the restore completes.
  If you select **Flatten existing directory structure** and you do not select **Overwrite existing files**, the restore succeeds. The first file that is restored is present when the restore completes. To prevent this issue, do not select **Flatten existing directory structure** when restoring multiple files with the same name.

- The **Flatten existing directory structure** and **Append string to file names** options are only applicable to files. They are not available for directories.
- Multiple restore jobs to the same VM are not supported. The user must start another job as needed for that VM once the first restore job for that VM has completed.
- If a backup and a restore occur simultaneously on the same VM, one or both jobs can have unexpected results. If a backup or a restore exits with a non-zero status code, one possible cause is simultaneous jobs occurring on the same VM.
- Cohesity does not recommend VMware agentless restore if a NetBackup client already exists on the target VM. The NetBackup administrator must use the agent-based restore in such cases.
- For the current list of guest operating systems that NetBackup supports for the target VM, see *Supported guest operating systems for VMware* in the following document:
  Support for NetBackup in virtual environments

Provide access to a credential for agentless single file recovery to a guest VM

A VMware administrator that wants to perform an agentless single file recovery to a guest VM may not have access to a guest VM's credentials. You can give a user access to a credential through an RBAC role. Either of the following methods allows the user to perform a recovery with a stored credential so they don't need to know the actual username and password for the VM.

See "Add a credential for a VMware guest VM" on page 244.

Note: This credential type is not for VMware servers. Configure those credentials on the **VMware servers** tab in **Workloads > VMware**.

You can give a user access to a credential in the following ways.

- Add a user to the Default VMware Administrator role. This RBAC role allows users to view all credentials and use any credential for recovery.
- Create a custom role that has access to a limited number of credentials. Then add users to that role.
  See "Create a custom role for agentless single file recovery to a guest VM, with a credential" on page 245.

Add a credential for a VMware guest VM

This type of credential lets you save the credentials for a guest VM in credential management. You can give a VMware administrator access to this credential. This user can then perform an agentless single-file recovery to a guest VM with the saved credential. They do not need to know the actual username and password for the VM.

Note: This credential type is not for VMware servers. Configure those credentials on the **VMware servers** tab in **Workloads > VMware**.

See "Add VMware servers" on page 66.

To add a credential for a VMware guest VM

1. On the left, click **Credential management**.
2. On the **Named credentials** tab, click **Add**.
3. Provide the following properties.
   - Credential name
   - Tag
   - Description
     For example, "This credential is used to recover to a VMware guest VM."
4. Click **Next**.
5. Select **VMware guest VM**.
6. Provide the credential details that are needed for authentication.
7. Click **Next**.
8. Add a role that you want to have access to the credential.
   - Click **Add**.
   - Select one of the following roles.
     The **Default VMware Administrator** role. This role has access to any and all credentials that are created.
   - Another role that has the necessary permissions to perform VMware single file recovery operations.
     Minimally the role should have the permissions **View** and **Assign credentials**.
9. Click **Next** and follow the prompts to complete the wizard.

Create a custom role for agentless single file recovery to a guest VM, with a credential

A custom role can allow a VM administrator to perform an agentless single file recovery to a guest VM, with a stored credential. This way the user doesn't need to know the actual username and password for the VM.

Use this role if you do not want users to have the Default VMware Administrator role. Or, you do not want to give users access to all credentials.

To create a custom role for agentless single file recovery to a guest VM, with a credential

1. A credential must exist that contains the username and password for the guest VM.

   See "Add a credential for a VMware guest VM" on page 244.

   Contact your NetBackup administrator for assistance.
2. On the left, select **Security > RBAC** and click **Add**.
3. Select **Default VMware Administrator** and click **Next**.
4. Provide a **Role name** and a description.

   For example, include a description that the role allows users to perform a single file recovery to a particular guest VM.
5. Under **Credentials**, click **Edit**.
6. Clear the option **Apply permissions to new and existing credentials**.
7. Select the credentials that you want to add to the role. Then click **Assign**.

   Users with the role have access to each credential that you select.
8. Under **Users**, click **Edit**. Then add the users that you want to have this RBAC role.
9. When you are done configuring the role, click **Add role**.

Recover files and folders with VMware agentless restore

This type of restore requires that you provide the credentials for the guest VM. Or, that you have access to the **VMware guest VM** credential that is saved in NetBackup credential management. Contact your NetBackup administrator for details.

To restore VMware files and folders using agentless restore

1. Confirm that the target machine is powered on.
2. On the left, click **Workloads > VMware**.
3. Locate and click on the VM that contains the files and folders for restore.
4. Click the **Recovery points** tab. In the calendar view, click the date on which the backup occurred.

   The available images are listed in rows with the backup timestamp for each image.
5. On the image you want to recover from, click **Recover > Restore files and folders**.
6. On the **Add files** page, click **Add** and select the files and folders you want recover. Click **Next**

   If you do not see the correct directory structure, click **Switch to instant access**. Note that instant access must be supported for the recovery point. If you still do not see the expected files and folders, start over and select a different recovery point.

   See "Create an instant access VM" on page 185.

   After you switch to instant access, all selected files are removed and all recovery options are reset. A new recovery begins of files and folders using instant access. If you want to switch back to agentless single file recovery again, you need to cancel the recovery wizard and restart.
7. Select the agentless recovery type and specify the target machine to which you want the files and folders recovered.
8. Enter the credentials for the target guest VM. Or, click **Select existing credentials** to select the credential you want to use.
9. Click **Next**.
10. On the **Recovery options** page, specify additional recovery options for the restored files and folders. Click **Next**.

    NetBackup performs a pre-recovery check using the options you specified.
11. On the **Review** page, review the status of the pre-recovery check along with the options you selected for the recovery. Once you confirm that they are correct, click **Start recovery**.

About restricted restore mode

The restricted restore mode option is a form of VMware agentless restore for restricted environments such as Windows User Account Control (UAC). The user-selected files are first staged to the recovery host and then restored to the virtual machine. The recovery host must have sufficient space for staging.

The default staging location on the recovery host is *install_path*\VERITAS\NetBackup\var\temp\staging. NetBackup creates this directory with the correct permissions the first time it is accessed. You can change the staging location with the AGENTLESS_RHOST_STAGING_PATH registry setting on the recovery host. This REG_SZ registry key does not exist by default. It must be created in HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\NetBackup\CurrentVersion\Config.

If you change the staging location, Cohesity recommends that you let NetBackup create the staging directory. When you let NetBackup create the directory, the permissions are set correctly. For NetBackup to create the new staging directory, the immediate parent directory must exist. If you want the restore to use E:\recovery\staging, then E:\recovery must exist. If the E:\recovery directory does not exist, the restore fails.

If you create the directory yourself, the **SYSTEM**, the domain administrator, and the local administrator accounts must have **Full Control** permissions. Additionally, Access Control Lists inherited from the parent directory are not secure and must be disabled.

Restricted restore mode supports alternate location restores. You can configure the alternate location in the NetBackup web UI.

Limitations of restricted restore mode:

- Restricted restore mode is currently only supported on Windows. The recovery host must also be Windows.
- The file ownership of the restored files is set to the account that was used for the NetBackup backup operation.
- Restore of ACLs is not supported.
- Restricted restore mode does not support renaming of targets for soft links.
- Restricted restore mode creates new files where hard links had previously been used.
- Irregular files such as sparse files, device files, special files, and junction points are not supported.
- A supported version of VMware Tools must be running for the restore to succeed.

- File path length with the directory cannot exceed 260 characters.

Performance considerations

File transport through the required infrastructure for this restore method is significantly slower than VMware agentless restores. As a result of performance concerns, Cohesity recommends limiting the restore to fewer than 100 files and less than 1 GB of data.

Chapter 18

Restoring Individual files and folders from VMware backups

This chapter includes the following topics:

About restoring individual VMware files and folders

You can use the following methods to restore files and folders individually from a VMware backup. Individual file and folder recovery is supported from full and incremental backups, as long as the **Enable file recovery from VM backup** policy option is enabled.

Restore to a virtual machine on which NetBackup client software is installed.	The NetBackup web UI in NetBackup 10.3 and later supports this method. See "About VMware agentless restore" on page 240.
Restore to a virtual machine where the NetBackup client is *not* installed.	See "Recover files and folders with VMware agentless restore" on page 245.

Create an instant access VM.	See "Create an instant access VM" on page 185. See "Restore files and folders from a VM backup image" on page 187.
(Windows only) Restore to a virtual machine drive that is mapped to a host on which NetBackup client software is installed.	See "Setting up NetBackup Client Service for VMware restores to a Windows shared virtual machine drive" on page 253.
Restore to a *host* on which the NetBackup client software is installed (not to the virtual machine).	See the *NetBackup Backup, Archive, and Restore Getting Started Guide* on how to restore to different locations. Then manually copy the restored files to the virtual machine. (NetBackup does not perform this step.)

Limitations

The following limitations exist when you restore individual files and folders from a VMware backup image:

- Instant access does not support the restore of files and folders to a NetBackup client.

Restore individual files and folders

To restore individual files and folders

1. On the left, click **Workloads > VMware**.
2. Locate and click on the VM that contains the files and folders for restore.
3. Click the **Recovery points** tab, in the calendar view, click the date on which the backup occurred. The available images are listed in rows with the backup timestamp for each image.
4. On the image you want to recover from, click **Recover > Restore files and folders**.
5. On the **Add files** page, click **Add** and select the files and folders you want recovered. Click **Add** then click **Next**.
6. On the **Recovery target** page, click on the **Target machine** field. Select the **NetBackup client** recovery type and the target computer to which you want the files and folders recovered. Click **Select**. Select a restore target option. Click **Next**.
7. On the **Recovery options** page, specify additional recovery options for the restored files and folders. Click **Next**.
8. On the **Review** page, review the options you selected for the recovery. Once you confirm that they are correct, click **Start recovery** to initiate the recovery.

Recovery options for restore of VMware files

This topic describes the options for restoring individual folders and files from a VMware virtual machine backup.

Table 18-1 Options for individual file restore

Option	Description
RECOVERY options	Select from the following options.
Restore everything to original location	Restores the folders and files to the location where they resided when the backup occurred.
Restore everything to a different location	Restores the folders and files with their original hierarchy, but to a different location. Use the **Destination** field to enter the restore location. Click **Browse** to browse to the restore location.
Restore individual directories and files to different locations	Restores the folders and files to individually designated locations. To designate a restore destination for each source folder, click **Edit file paths**. To restore to a Windows mounted drive: Destinations must be entered as UNC path names that refer to shared drives on the virtual machine. For example, to restore the file `E:\folder1\file1` on virtual machine `vm1`, enter the following destination: `\\vm1\e$\folder1\file1` See "Setting up NetBackup Client Service for VMware restores to a Windows shared virtual machine drive" on page 253.
Allow overwrite of existing files	By default, this option is not selected to avoid overwriting a current file. Select this option to replace a file with the same name in the destination directory with the file you want to restore.
Restore directories without crossing mount points	By default, all file systems that are mounted in the selected directories are restored. Select this option to restore the selected directories without restoring all file systems that are mounted in those directories. **Note:** Mount points inside a backup image are always restored whether or not this option is selected.

Table 18-1 Options for individual file restore *(continued)*

Option	Description
Rename hard links	UNIX and Linux systems only. By default, hard link path names are restored exactly as they exist in the backup. Select this option to rename the hard link path names, if any exist. Cohesity recommends that you select this option in the following situations: ■ You restore system files to an alternate disk and not to the current system disk. ■ You use the alternate disk as the system disk with the original file paths. In this situation, Cohesity recommends that you select **Rename hard links**. Then, make sure that the option **Rename soft links** is not selected so that you can use the alternate disk and still have the correct file paths.
Rename soft links	UNIX and Linux systems only. By default, soft (symbolic) link path names are restored exactly as they exist in the backup. Select this option to rename the soft link path names, if any exist. Cohesity recommends that you do not select this option if you rename hard links.
Media server	You can use this option to select a media server that has access to the storage unit that contains the backup image. An example of such an environment is a Media Server Deduplication Pool (MSDP) with multiple media servers. **Note:** If the storage unit that contains the backup image is not shared with multiple media servers, this option is grayed out.
Job priority	Determines the restore job's priority for restore resources. A higher priority means that NetBackup assigns the first available drive to the first restore job with the highest priority. Enter a number (maximum 99999). The default for all restore jobs is 0, the lowest priority possible. Any restore job with a priority greater than zero has priority over the default setting.

See "About restoring individual VMware files and folders" on page 249.

Setting up NetBackup Client Service for VMware restores to a Windows shared virtual machine drive

To restore individual files to a Windows virtual machine that has a shared drive, note: the NetBackup Client Service must be logged on under an account that has Administrator privileges (not as the Local System account). An account with Administrator privileges lets NetBackup write to the directories on the virtual machine to which the data is restored.

If you try to restore files while the NetBackup Client Service is logged on as the Local System account, the restore fails.

To log on the NetBackup Client Service as Administrator

1. In Windows Services on the VMware recovery host, double-click the NetBackup Client Service.
2. Check the **Log On** tab: if the service is not logged on under an account that has Administrator privileges, stop the service.
3. Change the logon to the Administrator account, or to an account that has Administrator privileges.

   The account must have Administrator privileges in the domain in which both the virtual machine and the VMware backup host reside.
4. Restart the service.
5. Retry the restore.

Chapter 19

Using NetBackup to back up Cloud Director environments

This chapter includes the following topics:

About NetBackup for vCloud Director

NetBackup can back up VMware vCloud Director environments and restore virtual machines into vCloud Director.

Note: The NetBackup for VMware restrictions also apply to vCloud Director objects. VMware restrictions also may apply; see your VMware documention.

Table 19-1 describes the configuration requirements for backup of vCloud Director.

See "NetBackup for VMware: notes and restrictions" on page 35.

Table 19-1 Configuration for backup of vCloud Director virtual machines

Task	Description
Enter NetBackup credentials for the vCloud Director server and for its vCenter servers.	See "Add VMware servers" on page 66.
Configure the policy **Clients** tab	Select the following: ■ **Select automatically through VMware Intelligent Policy query** ■ **Enable VMware Cloud Director integration** With this option, the policy selects only vCloud-managed virtual machines for backup: it skips the virtual machines that are not in vCloud. NetBackup collects information on the vCloud environment, such as its organizations, virtual datacenters, and vApps. NetBackup also retrieves information about a vApp for later restore of the vApp and its virtual machines. **Note:** **Enable VMware Cloud Director integration** makes several vCloud keywords available in the policy Query Builder Field, for rule-based selection of virtual machines. If **Enable VMware Cloud Director integration** is not selected, NetBackup cannot use the keywords to locate virtual machines in vCloud Director and the backup fails. **Note:** The browse icon (next to the **Query Builder** fields) may list non-vCloud objects. If you select an object that is not in vCloud Director, it is excluded from the backup. See "Configure a VMware policy" on page 87. See "Notes on creating a NetBackup policy for vCloud" on page 255.

Notes on creating a NetBackup policy for vCloud

When you create a backup policy for vCloud virtual machines, note the following:

- The configuration requirements are described in the following topic:
  See "About NetBackup for vCloud Director" on page 254.
- To back up all the existing vApp templates, use the **vCDIsvAppTemplate** keyword in the Query Builder (**vCDIsvAppTemplate Equal TRUE**).
- To back up specific vApp templates, use the **vCDvApp** keyword in the Query Builder with appropriate operator and values to select the particular templates.
- To allow a policy to back up virtual machines from multiple vCloud Director organizations: On the **VMware** tab of the policy locate and expand the **VMware advanced attributes**, then enable **Multiple organizations per policy**.

Notes on restoring virtual machines into vCloud Director

To restore a virtual machine into vCloud Director, note the following:

- The Backup, Archive, and Restore interface allows the restore of one virtual machine at a time.
  You can use the `nbrestorevm` command to restore multiple virtual machines.
- When you back up a virtual machine in vCloud Director, use the **VMware display name** setting for **Primary VM identifier** on the **VMware** tab. Use of the **VM BIOS UUID** setting is not recommended.
- To restore a virtual machine into vCloud Director, the virtual machine must have been in vCloud Director when it was backed up.
  To restore a virtual machine into vCloud Director, the backup policy must have been configured as described in the following topic:
  See "About NetBackup for vCloud Director" on page 254.
- For vCloud Director restores, you can use the following configuration setting to have NetBackup automatically delete the VM left at the vCenter on import failure:
  `DELETE_VM_ON_IMPORT_FAILURE = 1`
  Enter this setting in `bp.conf` or the registry on the primary server.
- When a vApp is restored, vCloud Director resets the vApp's expiration date. For example: Assume the original vApp was created on the first day of the month and was set to expire in 30 days. If the vApp is restored 15 days before its expiration (on the 15th), vCloud resets the vApp to expire in 30 days from the 15th.
  The VM administrator can reset the expiration date to its original date.
- You cannot restore a virtual machine into an existing vApp template. VMware sets this restriction.
- After you restore a vApp template, the template cannot be changed and no further virtual machines can be added to it (a VMware restriction). To restore a vApp template that is to contain multiple VMs, you must restore all but one of the VMs separately into a non-template vApp. Then restore the last virtual machine by means of the **Capture vApp as a template in catalog** option.
  See "Restore a vApp template that has multiple virtual machines" on page 262.
- You can restore the vCloud Director virtual machine into vSphere instead of vCloud Director. On the **Recovery target** page, select **vSphere.**
- vCloud organization networks are not displayed for restore; only vSphere networks are displayed.

- To restore into an existing vCloud Director vApp with the **Capture vApp as a template in catalog** option, the vApp must be turned off.
- vCloud backup images cannot be restored by means of the NetBackup vSphere Client (HTML5) plug-in. This type of restore can be performed by means of the Backup, Archive, and Restore interface.
- To ensure that any VM guest customizations are restored into vCloud Director, you must set a NetBackup parameter. The parameter value specifies a wait period in seconds so that the guest customizations can be restored successfully. (The VMware API requires that the VMware Tools are installed and running, but the state of the VMware Tools cannot be identified after the restore. Therefore, we wait the specified amount of time so that the VMware Tools are running in the initial restore environment.)
  See "Ensuring that guest customizations can be restored in vCloud Director" on page 354.
- If a VMware Cloud Director Organization Virtual datacenter (VDC) is configured with `Fast Provisioning` enabled, it does not allow different storage policies to be associated with a VM's home directory, virtual disks or both.
- If a VMware Cloud Director Organization Virtual datacenter (VDC) is configured with `Fast Provisioning` enabled and different storage policies are configured for a VMware Cloud Director VM's home, its virtual disks or both then, the storage policies are not applied and the restore job status is set to 1.
- If `Fast Provisioning` is disabled then a VMware Cloud Director VM allows for different storage policies to be associated with a VM's home directory, virtual disks or both.

Recover VMware Cloud Director virtual machines

You can only recover a virtual machine (VM) to a VMware Cloud Director if the VM was backed up from a VMware Cloud Director.

To recover a VMware Cloud Director VM

1. On the left, select **Workloads > VMware >** and select the virtual machine to recover.
2. Click the **Recovery points** tab. In the calendar view on the left, select the date on which the backup occurred.
3. For the image that you want to recover, select one of the following image recovery options:
   - **Recover**

Recover from the default copy of the backup image. This option is displayed if only one copy exists.

- **Recover from default copy**
  Recover from the default copy of the backup image. This option is displayed if more than one copy exists.
- ***nn* copies**
  Recover from the default copy or a different copy of the backup image. NetBackup allows up to ten copies of the same backup image. All available copies are displayed when you select this option. For each copy, the **Storage** name, **Storage server**, and the **Storage server type** are displayed.

4 Choose the recovery type **Restore virtual machine**.

5 On the **Recovery target** page, select to restore the VM to either **VMware Cloud Director** or **vSphere**.

- If you select **vSphere**, refer to the following information:
  See "Recover a full VMware virtual machine" on page 231.
- If you select **VMware Cloud Director**, continue with this procedure.

6 On the **Recovery target** page, specify the VMware Cloud Director and vSphere recovery destination information.

- The default values shown restore the VM back to its original location.
- If you change any of the VMware Cloud Director recovery destination information, you must update the **vSphere** recovery destination information.
- If you accept the default VMware Cloud Director recovery destination information, you can change the **vSphere** recovery destination information if necessary.

See "Recovery target" on page 259.

Click **Next**.

7 On the **vApp options** screen, specify the vApp information.

- To restore to an existing vApp, browse the list of vApps or enter the name of a vApp that exists.
- To restore to a new vApp, enter the name of the new vApp.
- The **Status** shows **New** if the vApp does not exist in VMware Cloud Director. A new vApp is created.

See "vApp options" on page 260.

Click **Next**.

8 For the **Recovery options** page, specify any recovery options for your restore and click **Next**.

9 The **Review** screen summarizes the selections made. A pre-recovery check attempts to determine if there are issues with any of the selected options. You can override any errors shown, however, the recovery can fail if errors are not addressed.

Recovery target

Table 19-2 Recovery target options

Option	Description
Restore to	**VMware Cloud Director** Select this option to restore a virtual machine into vCloud Director. **vSphere** Select this option to restore the virtual machine to the original location or to an alternate location.
VMware Cloud Director recovery destination	These options display when you choose to restore with the option **VMware Cloud Director**. **Organization vCD** The organization virtual datacenter. NetBackup displays the name of the **vCloud server** and the **Organization**.

Table 19-2 Recovery target options *(continued)*

Option	Description
vSphere recovery destination	**Display name** Specifies the VMware display name for the restored virtual machine. The default is the display name that the virtual machine had when it was backed up. The display name must be unique for the vCenter Server where the virtual machine is restored. **Note:** If a virtual machine with this display name already exists at this location (or at the original location), you are prompted to click overwrite the existing virtual machine. You cannot restore the virtual machine if the result is two virtual machines with the same display name on the same vCenter server. **ESXi server or cluster** Specifies the ESX server or cluster on which the restored virtual machine is to reside. To use the original ESX server or cluster (the default), verify that the original ESX server or cluster still exists. **Resource pool or vApp** Use this option to have the restored virtual machine assigned to either a VMware resource pool or to a vApp. Resource pools manage the host's CPU and memory. vApps are logical containers for virtual machines, and also share some functionality with virtual machines. **Datastore or datastore cluster** Specifies the VMware datastore or datastore cluster that contains the virtual machine configuration files. This datastore (sometimes called the vmx directory) contains the configuration files that describe the virtual machine, such as `*.vmx` files. Active snapshots of vmdk files are also stored on this datastore. **Note:** The Datastore field shows the name of the datastore that contained the virtual machine data when the virtual machine was backed up. Even if the datastore was in a datastore cluster, the field shows the name of the datastore, not the datastore cluster. When the virtual machine is restored, NetBackup determines how the datastore is currently configured (in a cluster or not) and configures the virtual machine accordingly.

vApp options

Table 19-3 Options to restore to a vApp

Field	Description
vApp name	Select the name of the vApp for the restore. This option defaults to the original vApp that was recorded in the backup.

Table 19-3 Options to restore to a vApp *(continued)*

Field	Description
Overwrite the existing virtual machine if it exists	Overwrites any virtual machine if it exists. If the vApp has no VMs, this option is disabled.
Remove exising vApp (or vApp template) and recreate it	Removes the existing vApp or vApp template and recreates the vApp or vApp template.
Capture vApp as a template in catalog	This option is available when you restore to an existing vApp. ■ Leave this option disabled to create a new vApp for the restore. The virtual machine is restored into a new vApp template. The name of the new template is the same as the vApp that was specified in the vApp Name field. ■ Enable this option to copy the vApp that is specified in the **vApp Name** field into a new vApp template. It also copies all of the vApp's virtual machines into the same vApp template. Note that this operation may take a lot of time. By default, the source vApp for the copy is retained after the copy; you can have the source vApp removed after the copy. Select the **Remove vApp after capture** option. **Catalog** Select the catalog in which to place the vApp template. The organization determines the available catalogs. **vApp template name** Enter the name for the new vApp template. **Remove vApp after capture** Select this option to remove the source vApp after the vApp copy to the new vApp template completes. At the end of the copy, the new template and its virtual machines are retained. The vApp that was the source for the copy is deleted.

Recovery options

Allow overwrite of existing virtual machine
NetBackup deletes any VM with the same display name that exists at the destination, before the recovery starts. Note that, NetBackup deletes any VM with the same display name, it may not be the same VM, but another VM having the same display name.

For VMware Cloud Director recovery, this option is not displayed if you choose to restore to **vSphere** (not **VMware Cloud Director**).

Power on after recovery
Automatically turns on the VM when the recovery is complete.

Recovery host
Indicate the host that you want to use to perform the recovery. By default, the recovery host is the one that performed the backup.

Restore a vApp template that has multiple virtual machines

To restore a vApp template that has multiple virtual machines

1. Use the Backup, Archive, and Restore interface to restore all but one of the virtual machines into a non-template vApp.

   You can restore one virtual machine at a time. After you have restored the first virtual machine, restore the second virtual machine with **Restore into existing vApp** on the **Recovery vApp Options for vCloud Director** screen. Select the same vApp into which you restored the first virtual machine. Step through the restore screens to restore each virtual machine in this way, except for the last virtual machine.

   Note: Use the following steps to restore the last virtual machine and to copy all the restored virtual machines into a vApp template.

   ---

2. On the **Recovery Destination** screen, select **Alternate location in vCloud Director**.
3. On the **Recovery vApp Options for vCloud Director** screen, do the following:
   - Select **Restore into existing vApp**.
   - Select the vCloud server and the organization that includes the vApp into which you have restored the other virtual machines.
   - If necessary, browse for the vApp into which you have restored the other virtual machines.
   - Select **Capture vApp as a template in catalog**.
   - Select the catalog to contain the template vApp.
     Note: The organization determines the available catalogs.
   - Enter a name for the vApp template.
   - **Remove vApp after capture**: Deletes the non-template vApp into which you restored the other virtual machines at the beginning of this procedure. Use this option to free up space on the datastore after the restore is complete.
4. On the **Recovery Destination Options for vCloud Director** screen, select the last virtual machine that you want to restore into the template vApp.
5. On the **Virtual Machine Options** screen, select the appropriate options for the virtual machine and its disk provisioning.

6 On the **Network Connections** screen, select the network for the restored virtual machine.

7 On the **Perform Recovery** screen, run a pre-recovery check.

To begin the restore click **Start Recovery**.

NetBackup copies the current virtual machine and the previously restored virtual machines into a new vApp template. When the restore is complete, no further virtual machines can be added to the template vApp.

Reducing the time required for VM discovery in a large vCloud environment

NetBackup backup policies for vCloud Director use query rules to automatically search and filter the vCloud environment. By default, the query rules search all the vCloud Director (vCD) servers in your environment. If the environment contains many vCloud servers with many vApps, VM discovery may take a long time. You can speed up VM discovery by limiting the search to specific vCloud servers or vApps.

The following is an example of a policy Query Builder rule that searches all vCloud servers and all vApps:

```
vmware:/?filter=vCDvApp Contains "vapp1"
```

With the Query Builder in the NetBackup web UI, the following is an example that uses OData keywords:

```
vmware:/?filter=Displayname contains(vcdvApp, 'vapp1')
```

To limit the search to particular servers or vApps, insert an additional vCloud expression in the Query Builder rule as explained in the following procedure.

To reduce the time required for VM discovery in a vCloud environment

1 In the NetBackup web UI, open the vCloud Director policy.

2 On the **Clients** tab of the policy, make sure **Select automatically through VMware Intelligent Policy query** and **Enable Cloud Director integration** are selected.

3 In the **Query Builder**, click **Advanced mode**.

4 Create one or more rules to search for VMs in specific vCloud Director servers or vApps.

Each query rule must begin on its own line.

You can use the following types of rules:

- To search for VMs in a particular vCloud server

  `vmware://<vCloud_server>?filter=<filter>`

- To search for VMs in a particular vApp or vApp template

  `vmware:/vApp/vapp-<vApp_id>?filter=<filter>`

  `vmware:/vAppTemplate/vappTemplate-<vAppTemplate_id>?filter=<filter>`

- To search for VMs in a particular vApp or vApp template on a particular vCloud server

  `vmware://<vCloud_server>/vApp/vapp-<vApp_id>?filter=<filter>`

  `vmware://<vCloud_server>/vAppTemplate/vappTemplate-<vAppTemplate_id>?filter=<filter>`

Example Query Builder rules for searching specific vCloud servers or vApps

5 For two or more search rules, you must enable multiple organizations for the policy.

- In the policy **VMware** tab, locate and expand **VMware advanced attributes**.
- Enable **Multiple organizations per policy**.
- To use the `nbdiscover` command instead of the policy **Query Builder**, see the following topic:

  Examples of the nbdiscover command for searching specific vCloud servers or vApps

Example Query Builder rules for searching specific vCloud servers or vApps

In the NetBackup policy Query Builder, you can speed up discovery of vCloud VMs by using the following types of query rules:

- To search for VMs in a particular vCloud server (note the double forward slash):

  `vmware://<vCloud_server>?filter=<filter>`

  Example rule:

  `vmware://vCD1.acme.com?filter=vCDvApp Contains "vapp1"`

  Example using OData keywords in the NetBackup web UI:

  `vmware://vCD1.acme.com?filter=contains(vcdvApp, 'vapp1')`

  NetBackup searches for VMs only in the `vCD1.acme.com` server.

- To search for VMs in a particular vApp or vApp template (note the single forward slash):

  `vmware:/vApp/vapp-<vApp_id>?filter=<filter>`

  `vmware:/vAppTemplate/vappTemplate-<vAppTemplate_id>?filter=<filter>`

The vApp_id or vAppTemplate_id is the identifier on the end of the vCloud vApp `href`. You can use a vCloud Director REST API query to find the identifier.
For example, the following is a REST API query for a vApp that is named `acmvappvm7`:

```
https://acmvm5.acme.com/api/query?type=adminVApp&filter=
name==acmvappvm7
```

The following example is an excerpt from the API query Response:

```
href="https://acmvm5.acme.com/api/vApp/vapp-afaafb99
-228c-4838-ad07-5bf3aa649d42"
```

In this example, the vApp identifier for vApp `acmvappvm7` is `afaafb99-228c-4838-ad07-5bf3aa649d42`. You can use this identifier in a NetBackup Query Builder rule as follows:

```
vmware:/vApp/vapp-afaafb99-228c-4838-ad07-5bf3aa649d42
?filter=Displayname Contains "prod"
```

Where `Displayname Contains "prod"` is an example filter for the rule.
Example using OData keywords in the NetBackup web UI:
`vmware:/vApp/vapp-afaafb99-228c-4838-ad07-5bf3aa649d42?filter=contains(displayName, 'prod')`

- To search for VMs in a particular vApp on a particular vCloud Director server:
  `vmware://<vCloud_server>/vApp/vapp-<vApp_id>?filter=<filter>`
  Example rule:

  ```
  vmware://vCD1.acme.com/vApp/vapp-4c0d9722-80a4-4f19-b636-72ebf48e4e71
  ?filter=Displayname Contains "prod"
  ```

  Example using OData keywords in the NetBackup web UI:
  `vmware://vCD1.acme.com/vApp/vapp-4c0d9722-80a4-4f19-b636-72ebf48e4e71?filter=contains(displayName, 'prod')`
- To search additional vCloud Director servers or vApps from the same backup policy, include additional query rules in the Query Builder.

Note: To enter multiple rules in the **Query Builder**, you must be in **Advanced mode**.

Note: Start each rule on its own line.

Example of two rules in the Query Builder:

```
vmware://vCD1.acme.com/vApp/vapp-4c0d9722-80a4-4f19-b636-72ebf48e4e71
?filter=Displayname Contains "prod"
vmware://vCD2.acme.com/vApp/vapp-5c0c9833-80a4-4f19-b636-72ebf48e4e63
?filter=Displayname Contains "prod"
```

Examples using OData keywords in the NetBackup web UI:

```
vmware://vCD1.acme.com/vApp/vapp-4c0d9722-80a4-4f19-b636-72ebf48e4e71
?filter=contains(displayName, 'prod')
vmware://vCD2.acme.com/vApp/vapp-5c0c9833-80a4-4f19-b636-72ebf48e4e63
?filter=contains(displayName, 'prod')
```

Examples of the nbdiscover command for searching specific vCloud servers or vApps

Use the following `nbdiscover` command format to search for VMs in specific vCloud servers or vCloud vApps:

```
nbdiscover "<vCloud_query>" -job_info "snaparg=enable_vCloud=1"
```

Example 1. Search for VMs in vCloud server `vCD1.acme.com` only:

```
nbdiscover -noxmloutput "vmware://vCD1.acme.com?filter=DisplayName
Contains 'prod1'" -job_info "snaparg=enable_vCloud=1"
```

The `-noxmloutput` option displays one VM per line.

Example 2. Search for VMs in a vCloud vApp that has the following vApp identifier:

```
4c0d9722-80a4-4f19-b636-72ebf48e4e71
```

```
nbdiscover -noxmloutput "vmware:/vApp/vapp-4c0d9722-80a4-4f19-b636
-72ebf48e4e71?filter=DisplayName Contains 'prod1'" -job_info
"snaparg=enable_vCloud=1"
```

The following topic provides assistance in finding the vApp identifier:

See the section called "Example Query Builder rules for searching specific vCloud servers or vApps" on page 264.

Example 3. Use two query rules to search for VMs prod1 and prod2 in a vCloud vApp:

```
nbdiscover -noxmloutput "vmware:/vApp/vapp-4c0d9722-80a4-4f19-b636
-72ebf48e4e71?filter=DisplayName Contains 'prod1'" "vmware:/vApp/
vapp-4c0d9722-80a4-4f19-b636-72ebf48e4e71?filter=DisplayName Contains
 'prod2'" -job_info "snaparg=enable_vCloud=1,multi_org=1"
```

Note the two rules, each enclosed with double quotes and separated by a space, and the `multi_org=1` option.

Chapter 20

Restore virtual machines with Instant Recovery

This chapter includes the following topics:

About Instant Recovery for VMware

NetBackup can recover a virtual machine almost instantly, without waiting to transfer the virtual machine's data from the backup. NetBackup starts the virtual machine

directly from the backup image and makes it accessible to users on the target ESX host immediately. You can copy files (including vmdk files) without restoring the entire virtual machine. To restore the virtual machine, use VMware Storage vMotion to migrate the virtual machine data files from the backup image to the ESX host.

Some examples of instant recovery:

- Access and restore individual files and folders from any type of OS and then delete the virtual machine. (Note for Windows or Linux: Instead of instant recovery, you can use the policy option **Enable file recovery from VM backup** and restore individual files with the Backup, Archive, and Restore interface.)
- Test a patch on a restored virtual machine before you apply the patch to production systems.
- Troubleshoot a virtual machine or host, such as when the production ESX host is down. You can start the virtual machine from its backup and use it until the production system is back online.
- Permanently recover the virtual machine by means of Storage vMotion.
- Verify the backup image.
- Copy a vmdk file and then delete the virtual machine.
- Verify an application.

In any case, the virtual machine is started directly from the backup image and is available in seconds or minutes. The startup time depends on the network speed and storage speed, not on the size of the virtual machine.

Table 20-1 describes the steps in a virtual machine instant recovery.

Table 20-1 How Instant Recovery for VMware works

Sequence	Actions
Step 1	Run the nbrestorevm command* to access the virtual machine from its backup image. The NetBackup File System Service (NBFSD) on the media server accesses the backup image file system and mounts the image as an NFS datastore. The datastore becomes accessible to the ESX host where the virtual machine is to be restored. On the same command, select a temporary datastore that is accessible to the ESX host.
Step 2	NetBackup creates a virtual machine on the ESX host and configures the virtual machine with write access to a temporary (local) datastore.
Step 3	NetBackup creates a snapshot of the virtual machine. Any new write requests in the virtual machine use the temporary datastore. The virtual machine uses the NFS datastore as read only.

Table 20-1 How Instant Recovery for VMware works *(continued)*

Sequence	Actions
Step 4	NetBackup starts up the virtual machine on the ESX host.
Step 5	To keep the restored VM: Use Storage vMotion to copy the virtual machine data from the NFS datastore to the temporary datastore.
Step 6	When vMotion is complete, use nbrestorevm to unmount the NFS datastore.

*NetBackup provides a command-line interface for instant recovery of virtual machines (nbrestorevm). A graphical user interface (the Instant Recovery Wizard) is available in the NetBackup vSphere Client (HTML5) plug-in. For details, see the NetBackup Plug-in for VMware vSphere Client (HTML5) Guide.

Task overview for Instant Recovery for VMware

Table 20-2 describes the tasks for Instant Recovery for VMware.

Table 20-2 Instant Recovery tasks

Step	Description	Reference topic
Step 1	Review the performance recommendations	See "Performance recommendations for Instant Recovery for VMware" on page 271.
Step 2	Review the notes and requirements	See "Requirements for Instant Recovery for VMware" on page 271. See "Notes on Instant Recovery for VMware" on page 272.
Step 3	Restart the Client for NFS service on the restore host	See "Restarting the Client for NFS service on a Windows restore host" on page 274.
Step 4	Review the Instant Recovery options on the nbrestorevm command	See "Instant Recovery options on the nbrestorevm command" on page 275.
Step 5	Use the nbrestorevm command to perform Instant Recovery	See "Restoring a virtual machine with Instant Recovery for VMware" on page 276. See "Restoring individual files with Instant Recovery for VMware while the current virtual machine is running" on page 283.

Performance recommendations for Instant Recovery for VMware

After Instant Recovery, the virtual machine is on an NFS-attached datastore that is presented by the NetBackup media server. The performance of the virtual machine and of Storage vMotion depends on the following: the network speed and latency between the ESXi host and the media server, and the speed of NetBackup storage that the backup is recovered from.

It is recommended the following:

- A SAN connection from the NetBackup media server to its disk storage unit.
- For Fibre Channel SAN, a minimum speed of 4 gigabits per second.
- For iSCSI SAN, a minimum speed of 1 gigabit per second.
- When you use Storage vMotion to migrate a restored virtual machine, migrate one virtual machine at a time per media server. The migration may be slow if you simultaneously migrate multiple virtual machines per media server.
- For disaster recovery testing, it is recommended that you restore no more than three or four virtual machines per media server. The number to restore depends on the I/O load on the media server. It is recommended restoring each VM one-by-one, not simultaneously.

Note: For large-scale recovery of multiple virtual machines, use the virtual machine restore feature in the Backup, Archive, and Restore interface. Do not use Instant Recovery for VMware.

Requirements for Instant Recovery for VMware

For virtual machine instant recovery, your environment must meet the following requirements:

- The virtual machine to restore must have been backed up from a VMware policy.
- The target ESX server for the restore must be at vSphere 5.0 or later.
- The restore host can be on Windows or Linux.

Note: For the VMware virtual machines that have non-ASCII characters in their paths, NetBackup does not support Instant Recovery using Windows restore hosts and media servers. You must use a Linux restore host and a Linux media server for Instant Recovery of such virtual machines.

The requirements and limitations for non-ASCII character support are described in a different topic.

See "NetBackup for VMware: notes and restrictions" on page 35.

- For a restore host that is separate from the NetBackup primary server or media server: You must add the restore host to the list of servers that can access the primary server.
  In the NetBackup web UI, select **Hosts > Host properties**. Select the primary server. Select **Connect**. Select **Servers**. On the **Additional servers** tab, select **Add** to add the restore host
- The NFS Client service must be enabled on the ESXi host.
- The Network File System (NFS) must be installed on the Linux media server and restore host and the `portmap` service must be active.
  For information about how to install NFS, see the media server host operating system documentation.
- The Services for Network File System (NFS) must be installed on the Windows media server and restore host.
  See "About configuring services for NFS on Windows 2012 or 2016 (NetBackup for VMware)" on page 361.
  The NetBackup media server platform must support Granular Recovery Technology. See the NetBackup Software Compatibility List (SCL).
- The Client for NFS service may have to be restarted on a NetBackup Windows restore host.
  See "Restarting the Client for NFS service on a Windows restore host" on page 274.
- The media server must use IPv4 or have a dual stack configuration if the vCenter server has a dual stack configuration.
- NetBackup requires logon credentials for the vCenter server and the restore host.
  See "Add VMware servers" on page 66.

Notes on Instant Recovery for VMware

Note the following about instant recovery of VMware virtual machines:

- Supports the following storage unit types (disk only): BasicDisk, AdvancedDisk, Media Server Deduplication Pool (MSDP), and qualified third-party OpenStorage devices.
  Note: Snapshot-only backups are not supported.
- Does not support a virtual machine that had the disks that were excluded from the backup. The policy **Virtual disk selection** option must have been set to include all disks.
- Does not support a virtual machine that has a disk in raw device mapping mode (RDM) or that has a disk in Persistent mode.
- Supports the following policy schedule types: Full backups, and the incremental backups that include the **Use Accelerator** option with a disk-based storage unit. Incrementals without the **Use Accelerator** policy option are not supported.
- Does not support virtual machine templates.
- If the virtual machine contains an IDE drive, the restored virtual machine may not start. This issue is not unique to instant recovery.
  See "VMware virtual machine does not restart after restore" on page 347.
- To avoid host name or IP address conflicts between the current virtual machine and the virtual machine version you want to restore: Shut down the virtual machine in your production environment before you start the recovery. Then change the display name of the current virtual machine, or use the `-R` option on nbrestorevm to rename the restored virtual machine.
- For a virtual machine that is running under a high load, migration of the virtual machine may take longer than expected. For this reason, NetBackup changes the virtual machine's `fsr.maxSwitchoverSeconds` property to 900.
  For example, this increase may be necessary when the virtual machine is restored from a deduplication storage unit.
  The following VMware Knowledge Base article contains more information on the `fsr.maxSwitchoverSeconds` property:
  Using Storage vMotion to migrate a virtual machine with many disks timeout
- Note the following about the virtual machine's datastore name:
  - If the name of the datastore includes spaces, the name should be enclosed in double quotes ("").
  - A virtual machine restore may fail if the name of the datastore (that was used at the time of the backup) ended with a period.
    The following tech note contains additional information.
    https://www.veritas.com/docs/100028139

- Instant recovery cannot restore a vCloud virtual machine into vCloud. The virtual machine is restored into vSphere. You can copy or import the restored virtual machine into vCloud by means of the **Copy** option in vCloud.
  Note that the vCloud **Move** option does not work with a virtual machine that runs from a NetBackup datastore.
- Storage lifecycle policies (SLPs) can use Auto Image Replication to replicate a virtual machine backup image to another NetBackup domain. To restore the virtual machine from the replicated image, you must include the -vmproxy option on the nbrestorevm command. Use the -vmproxy option to specify the backup host (access host) that is in the domain where the virtual machine was replicated. Without the -vmproxy option, nbrestorevm defaults to the backup host in the original domain and the restore fails.
- Supports recovery of virtual machines containing independent disks. The independent disks that are associated with the virtual machine are recovered to the virtual machine working directory on the temporary datastore. This functionality requires a NetBackup 8.3 or later recovery host.
- vMotion-derived restore (copy) of an Instant Recovery or Instance Access VM does not reinstate storage policies.

Restarting the Client for NFS service on a Windows restore host

It may be necessary to stop and restart the NFS Client service. If you use the Microsoft services snap-in (`Services.msc`) to restart it, the service does not start until you restart the server.

To restart the Client for NFS service without a server restart

- From the Windows command prompt, run the following commands:

```
net stop nfsclnt
net stop nfsrdr
net start nfsrdr
net start nfsclnt
```

  The Client for NFS service should restart without a restart of the server.

Instant Recovery options on the nbrestorevm command

NetBackup provides a command-line interface for instant recovery of virtual machines: the `nbrestorevm` command.The following is a list of the available `nbrestorevm` options for performing Instant Recovery of a VMware virtual machine.

Note: Although the nbrestorevm command has additional options, only the options that are described in this topic apply to Instant Recovery.

To initiate Instant Recovery (activate the virtual machine)

Options without brackets are required.

```
nbrestorevm -vmw -ir_activate -C vm_client
      -temp_location temp_location_for_writes
      [-S primary_server] [-vmpo] [-vmInstanceId] [-vmsn] [-vmst]
      [-vmserver vm_server] [-vmproxy vm_proxy] [-vmkeephv] [-vmid]
      [-vmnewdiskuuid] [-s mm/dd/yyyy [HH:MM:SS]]
      [-e mm/dd/yyyy [HH:MM:SS]]
      [-R absolute_path_to_rename_file]
      [-disk_media_server media_server]
```

Note: Only `-vmw`, `-ir_activate`, `-C`, and `-temp_location` are required. If the other options are not specified, NetBackup automatically supplies values for those options from the backup. In most cases, if you do not restore the virtual machine to a different location, you can omit the bracketed options.

To list details about the activated virtual machine

```
nbrestorevm -ir_listvm
```

To deactivate or delete the virtual machine

```
nbrestorevm -ir_deactivate instant_recovery_identifier [-force]
```

To complete the VM instant recovery job after the data is migrated

```
nbrestorevm -ir_done instant_recovery_identifier
```

To reactivate a virtual machine that was interrupted during recovery

```
nbrestorevm  -ir_reactivate instant_recovery_identifier [-force]
nbrestorevm -ir_reactivate_all -vmhost vm_host -media_server
      media_server_activate_vm [-force]
```

The command options are described in the *NetBackup Commands Reference Guide*.

Restoring a virtual machine with Instant Recovery for VMware

You can use this procedure to do either of the following:

- Copy files from a virtual machine backup.
- Restore the full virtual machine.

In either case, you can restore the virtual machine to its original location or to an alternate location.

Note: To avoid host name or IP address conflicts, shut down the current virtual machine in your production environment before you start instant recovery.

To copy files while the current virtual machine is running, use a different procedure:

See "Restoring individual files with Instant Recovery for VMware while the current virtual machine is running" on page 283.

Table 20-3 Basic steps for VMware instant recovery

Type of recovery	Steps
Copy files or troubleshoot an issue, then delete the restored virtual machine	Basic steps are these: ■ Restore the VM: Use nbrestorevm with the -ir_activate option. ■ Copy files from the VM; or use the VM as a stand-in until the production host is back online. ■ Delete the VM and release the media server resources: Use nbrestorevm with the -ir_deactivate option. See the following procedure for command details.

Table 20-3 Basic steps for VMware instant recovery *(continued)*

Type of recovery	Steps
Restore and keep the virtual machine	Basics steps are these: ■ Restore the VM: Use nbrestorevm with the ir_activate option. ■ Transfer the virtual machine files to an ESX host: Use the **Migrate** option in vSphere Client. ■ Release the media server resources: Use nbrestorevm with the ir_done option. See the following procedure for command details.

See "Requirements for Instant Recovery for VMware" on page 271.

The detailed procedure follows.

To restore a virtual machine with instant recovery

1 On the primary server, media server, or restore host, enter the nbrestorevm command.

This command is in the following location:

UNIX, Linux: /usr/openv/netbackup/bin/

Windows: *install_path*\NetBackup\bin\

Enter the command as follows.

To restore the VM to its original location:

```
nbrestorevm -vmw -ir_activate -C virtual_machine -temp_location
temporary_datastore [-vmproxy VMware_access_host] -vmpo
```

To restore the VM to a different location:

```
nbrestorevm -vmw -ir_activate -C virtual_machine -temp_location
temporary_datastore [-vmserver vCenter_server] -R rename_file_path
[-vmproxy VMware_access_host] -vmpo
```

`-C virtual_machine` identifies the virtual machine by the name or ID that was set in the policy's **Primary VM identifier** attribute for the backup. On the `-C` option, specify the same type of identifier that was used in the policy: VM host name, VM display name, VM BIOS UUID, VM DNS name, or VM instance UUID.

The `-R` option provides the path to a file that contains directives for restore to a different location.

See "Restoring a virtual machine to a different location with Instant Recovery for VMware" on page 280.

See "Instant Recovery options on the nbrestorevm command" on page 275.

The nbrestorevm command mounts the virtual machine's backup image as an NFS datastore and makes the datastore accessible to the ESX host. It also creates the VM on the ESX host. It then creates a snapshot of the virtual machine.

Note: Storage lifecycle policies (SLPs) can use Auto Image Replication to replicate a virtual machine backup image to another NetBackup domain. To restore the virtual machine from the replicated image, you must include the -vmproxy option on the command. Use the -vmproxy option to specify the backup host (access host) that is in the domain where the virtual machine was replicated. Without the -vmproxy option, nbrestorevm defaults to the backup host in the original domain and the restore fails.

The following tasks appear in the vSphere Client interface. In this example, dbl1vm5 is the virtual machine to be restored.

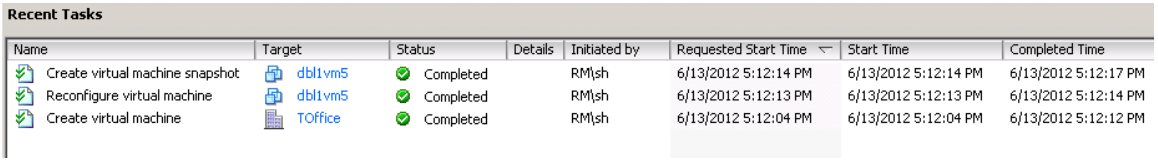

Recent Tasks

Name	Target	Status	Details	Initiated by	Requested Start Time	Start Time	Completed Time
Create virtual machine snapshot	dbl1vm5	Completed		RM\sh	6/13/2012 5:12:14 PM	6/13/2012 5:12:14 PM	6/13/2012 5:12:17 PM
Reconfigure virtual machine	dbl1vm5	Completed		RM\sh	6/13/2012 5:12:13 PM	6/13/2012 5:12:13 PM	6/13/2012 5:12:14 PM
Create virtual machine	TOffice	Completed		RM\sh	6/13/2012 5:12:04 PM	6/13/2012 5:12:04 PM	6/13/2012 5:12:12 PM

Note the following:

- The nbrestorevm command creates a NetBackup job of type "VM Instant Recovery."
  See "Job types for Instant Recovery for VMware" on page 286.
- If you cancel the instant recovery job or stop all NetBackup services, the NetBackup NFS datastore is unmounted and its media server resources are released.

  Caution: The virtual machine is deleted from the ESX host.

  ---

2 In vSphere Client, turn on the virtual machine.

  If you included the -vmpo option on the nbrestorevm command, the virtual machine is already turned on.

3 Browse and copy the virtual machine files as needed.

To copy files while the current virtual machine is running, use a different procedure.

See "Restoring individual files with Instant Recovery for VMware while the current virtual machine is running" on page 283.

4 If you do not want to keep the restored virtual machine, enter the following:

```
nbrestorevm -ir_listvm
```

In the output, find the VM Instant Recovery ID for the restored VM.

To remove the VM from the ESX host:

```
nbrestorevm -ir_deactivate instant recovery ID [-force]
```

where *instant recovery ID* is the virtual machine's numeric identifier from the -ir_listvm output. -force is optional, to suppress confirmation prompts.

The VM is removed from the ESX host. If no other VM uses the NetBackup NFS datastore, NetBackup removes that datastore and releases its resources on the media server.

The following tasks appear in the vSphere Client interface. In this example, dbl1vm5 is the virtual machine to be removed and datastore_V is the temporary datastore that it used.

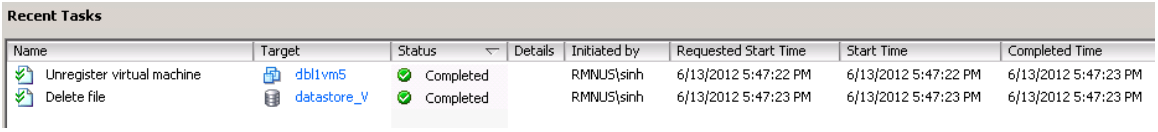

This step completes the VM Instant Recovery job. Skip the rest of this procedure.

Step 5 uses Storage vMotion to move the virtual machine to a production datastore. If vMotion is already in progress for this virtual machine, you should cancel the vMotion job before you enter `-ir_deactivate`. Otherwise, vMotion moves the virtual machine to a production datastore where `-ir_deactivate` cannot remove it.

5 To keep the restored virtual machine:

In vSphere Client, right-click on the restored virtual machine and select **Migrate**. Select the migration type and the destination.

Note: For the destination, select a permanent (production) location for the virtual machine. Do not select the temporary datastore that was used for the instant restore.

Storage vMotion transfers the virtual machine data files from the NetBackup NFS datastore to the datastore that you selected.

Note: You should migrate no more than one restored virtual machine at a time per media server.

6 After the migration is complete, use vSphere Client to merge or consolidate the virtual machine's redo log (or snapshot) files manually. See your VMware documentation for details.

When the migration to the production datastore is complete, use the following steps to unmount the NFS datastore and release its resources.

7 Enter the following:

```
nbrestorevm -ir_listvm
```

In the -ir_listvm output, find the VM Instant Recovery ID for the restored VM.

8 When the data migration is complete, enter the following:

```
nbrestorevm -ir_done instant recovery ID
```

where *instant recovery ID* is the virtual machine's numeric identifier from the -ir_listvm output.

The -ir_done option completes the VM Instant Recovery job. It also removes the NetBackup NFS datastore if no other VM uses it. When the datastore is removed, its resources are released on the media server.

Restoring a virtual machine to a different location with Instant Recovery for VMware

This topic explains how to use the nbrestorevm command to restore a VM to a different location.

The overall instant recovery procedure is available in another topic:

See “Restoring a virtual machine with Instant Recovery for VMware” on page 276.

To restore a VM to a different location

1 Find the path to a resource pool at the restore destination. (If you already know the full path, you can skip this step.)

Note: To restore to a different location, it is usually necessary to designate a different resource pool. Here is an example of a resource pool path:

```
/TechOffice/host/F2/pl9.acme.com/Resources
```

To find the path, enter the following on the primary server, media server, or restore host:

UNIX, Linux:

```
/usr/openv/netbackup/bin/bpVMreq <restore_host> 11 0 <ESXi_server>
<VMserver_or_vCenter_server>
```

Windows:

```
<install_path>\NetBackup\bin\bpVMreq.exe <restore_host> 11 0
<ESXi_server> <VMserver_or_vCenter_server>
```

Note: The numeric value `11 0` is required and must be entered as shown.

For example:

```
bpVMreq battleship.acme.com 11 0 ESXi_pl9.acme.com
vC_p9vm3.acme.com
```

Where `battleship.acme.com` is the restore host, `ESXi_pl9.acme.com` is the destination ESXi server, and `vC_p9vm3.acme.com` is the destination vCenter server.

This command generates a path to an XML file in a temporary location. The XML file lists all the available resource pools.

Here is an example of an XML file that `bpVMreq` creates:

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<ResourcePoolList><ResourcePool Name="Resources"
Path="/TechOffice/host/F2/pl9.acme.com/Resources"
Type="ResourcePool"><ResourcePoolList/>
</ResourcePool></ResourcePoolList>
```

In this example, the path to the resource pool is `/TechOffice/host/F2/pl9.acme.com/Resources`. Make a note of the path for use in the next step.

2 Create a text file with the following `change` entries.

Note: Each `change` entry helps to define the location for the restore. Each `change` line must end with a carriage return.

`change vmname to <new_virtual_machine_name>`	(The `change vmname` entry is optional)
`change esxhost to <new_ESXi_host>`	
`change resourcepool to <path_to_new_resouce_pool>`	
`change networkname to <new_network>`	(The `change networkname` entry is optional)

Enter each `change` line exactly as it appears in this list, except for the variable at the end (such as *new_virtual_machine_name*). Replace the variable with the new name. For example:

```
change esxhost to ESXi01.prod4.com
```

For the *new_resouce_pool*, use the path that was obtained in the first step of this procedure. For example:

```
change resourcepool to /TechOffice/host/F2/pl9.acme.com/Resources
```

This text file is called the `-R` rename file, and is used with the nbrestorevm command in the next step.

3 To restore the VM using the `-R` rename file, enter the `nbrestorevm` command with the `-R` option as follows.

Note: The `-R` option specifies the path to the text file (rename file).

- To restore to the same vCenter server but to a different ESXi host, enter the following:

  ```
  nbrestorevm -vmw -ir_activate -C <virtual_machine>
  -temp_location <temporary_datastore> -R <rename_file_path>
  ```

- To restore to a different vCenter server and a different ESXi host, enter the following:

  ```
  nbrestorevm -vmw -ir_activate -C <virtual_machine>
  -temp_location <temporary_datastore> -R <rename_file_path>
  -vmserver <vCenter_server>
  ```

To restore a VM after Auto Image Replication (AIR) to a disaster recovery (DR) site: you must also include the -vmproxy option on the nbrestorevm command to specify the restore host at the DR site.

Restoring individual files with Instant Recovery for VMware while the current virtual machine is running

You can use instant recovery to restore files individually from a virtual machine backup. You can restore the virtual machine from its backup image and mount it on a private network (such as a sandbox). This approach avoids the potential for network conflicts with the virtual machine in your production environment. Another virtual machine on a public network can be used as an intermediary, to copy the files from the virtual machine on the private network.

Note: This procedure lets you restore files into a running VM. You do not need to shut down the current virtual machine in your production environment before you start this procedure.

Before you start this procedure, you need an intermediary virtual machine that has a network connection to the public network or production network. In this procedure you connect the intermediary to the private network where the restored virtual machine is to be mounted.

At the end of the procedure, you can copy files from the restored virtual machine to the intermediary virtual machine. Then the virtual machines on the public network can access the files on the intermediary.

To restore individual files using instant recovery

1. Use vSphere Client to log on to the vCenter server.

   You must use a logon that allows access to the files that you want to recover.

2. Create a vSphere standard switch.

   This switch is for access to the ESX host from the sandbox or private network where the VM is to be activated from its backup.

   Note: The switch is for internal communication within the ESX host only.

   For example, in vSphere Client 5:

   - Select the ESX host for communication between the restored virtual machine and the intermediary virtual machine.
   - On the **Configuration** tab, in the **Hardware** pane, click **Networking**.
   - Click **Add Networking**.

- Select **Virtual Machine** as the connection type.
- Select **Create a vSphere standard switch**.
- For **Port Group Properties**, **Network Label**, enter a name for the internal switch (such as NB or NetBackup).
- Click **Finish**.

3 On the intermediary virtual machine, add a network card (NIC) to be connected to the vSphere standard switch.

Use this connection to retrieve files from the restored virtual machine that is to be mounted on the private network.

Note: This intermediary virtual machine must already have a network connection to the public network or production network.

For example, in vSphere Client 5:

- Select the intermediary virtual machine.
- On the **Summary** tab, click **Edit Settings**.
- Click **Add**.
- Select **Ethernet Adapter**.
- For the **Network label**, select the private network that is created in step 2.
- Click **Finish**.

4 If the intermediary's guest OS does not automatically assign an IP address for the private network after step 3, note: You must manually configure the IP address, default gateway, and subnet mask.

The intermediary should now be connected to both the public network and to the private network where the virtual machine is to be restored.

5 Use the nbrestorevm command to restore the virtual machine.

```
nbrestorevm -vmw -ir_activate -C virtual_machine -temp_location
temporary_datastore -R rename_file_path -vmsn
```

-vmsn specifies that no network is enabled for the virtual machine when it is activated from the backup image. Without the -vmsn option, network conflicts with the production virtual machine may occur.

The file that is designated by `-R rename_file_path` specifies a different display name or location for the restored virtual machine. You must change the virtual machine name or location to avoid conflicts with the current virtual machine in production. For example, to rename the virtual machine, the rename file can consist of the following entry (ending with a carriage return):

```
change vmname to acme_vm5
```

Note: The words `change vmname to` are literals, followed by the actual name to change to (such as `acme_vm5`).

See "Instant Recovery options on the nbrestorevm command" on page 275.

For other nbrestorevm options, see the man page or the *NetBackup Commands Reference Guide*.

6 Add a network card (NIC) to the restored virtual machine and connect the NIC to the vSphere standard switch from step 2.

7 Turn on the restored virtual machine.

8 If the guest OS does not automatically assign an IP address for the private network, configure the IP address, default gateway, and subnet mask.

9 Set up file sharing (such as through FTP, NFS, or CIFS) between the restored virtual machine and the intermediary virtual machine.

Then copy the files from the restored virtual machine to the intermediary virtual machine. The current virtual machine in production can access the files.

10 If you do not want to keep the restored virtual machine, enter the following:

```
nbrestorevm -ir_listvm
```

In the -ir_listvm output, find the VM Instant Recovery ID for the restored virtual machine.

To remove the restored virtual machine:

```
nbrestorevm -ir_deactivate instant recovery ID
```

where *instant recovery ID* is the virtual machine's numeric identifier from the -ir_listvm output.

Job types for Instant Recovery for VMware

Instant Recovery jobs appear as the following job types in the NetBackup Activity Monitor.

Table 20-4 VMware Instant Recovery job types in the Activity Monitor

Job type	Description

VM Instant Recovery	This job is the parent job for restoring a VM by means of Instant Recovery.
	To complete this job, you must enter one of the following:
	`nbrestorevm -ir_done instant recovery ID`
	`nbrestorevm -ir_deactivate instant recovery ID`
	For details on these commands:
	See "Restoring a virtual machine with Instant Recovery for VMware" on page 276.

Activate Instant Recovery	The parent VM Instant Recovery job starts an Activate Instant Recovery job to create the VM on the ESX host.
Stop Instant Recovery	This job runs when you use `nbrestorevm -ir_done` to remove the NetBackup NFS datastore and release its resources on the media server.
Deactivate Instant Recovery	This job runs when you use `nbrestorevm -ir_deactivate` to delete the restored VM from the ESX host.

Table 20-4 VMware Instant Recovery job types in the Activity Monitor *(continued)*

Job type	Description
Reactivate Instant Recovery	This job runs when you use `nbrestorevm` with the ir_reconfigure option to restart an interrupted virtual machine recovery.

Reactivating a restored virtual machine with Instant Recovery for VMware

If an interruption occurs during an instant recovery (such as a restart of the host or media server), the ESX connection to the media server may fail. In that case, it may be possible to re-establish the connection and return the virtual machine to the state it was in before the outage. Any transactions that occurred in the virtual machine before the outage are retained.

To reactivate a restored virtual machine

1. If only one VM had been restored to the ESX host, enter the following:

   ```
   nbrestorevm –ir_listvm
   ```

   Find the VM Instant Recovery ID for the restored VM in the -ir_listvm output. Then enter the following:

   ```
   nbrestorevm –ir_reactivate Instant Recovery ID [-force]
   ```

   where *instant recovery ID* is the virtual machine's numeric identifier from the -ir_listvm output. -force is an optional parameter to suppress confirmation prompts.

   The ir_reactivate option remounts the NetBackup NFS datastore. From the temporary datastore on the ESX host it registers the restored virtual machines on the ESX host.

2. If more than one VM had been restored to the ESX host:

   ```
   nbrestorevm –ir_reactivate_all -vmhost vm_host –media_server
   media_server [-force]
   ```

   Note: For multiple virtual machines, do not use the –ir_reactivate option. Use –ir_reactivate_all.

   The -vmhost option specifies the ESX host on which the virtual machines were mounted. The -media_server option specifies the media server on which the NFS datastores that contain the backup images were mounted. -force is an optional parameter to suppress confirmation prompts.

   The nbrestorevm -ir_reactivate_all command remounts the NetBackup NFS datastores on the media server and reactivates the virtual machines.

3. When the virtual machine is reactivated, you can copy its files or migrate its data to the ESX host.

   See “To restore a virtual machine with instant recovery” on page 277.

4. If Storage vMotion was migrating the virtual machine files when the outage occurred, restart the migration.

   In vSphere Client, right-click on the restored virtual machine and select **Migrate**.

Chapter 21

Protecting VMs using hardware snapshots and replication

This chapter includes the following topics:

About virtual machines and hardware snapshots

Hardware snapshot-based solution for VMware uses storage array snapshots for protecting VMware virtual machines. The benefits of using hardware snapshot are the reduced stun time for virtual machine. The VM snapshot is retained only for the duration of hardware snapshot.

This solution uses the NetBackup snapshot manager for performing hardware snapshots. For more information about NetBackup snapshot manager, refer to the *NetBackup™ Snapshot Manager for Data Center Administrator's guide.*

Deployment and architecture

Following is the deployment and architecture diagram of VMware hardware snapshot-based solution.

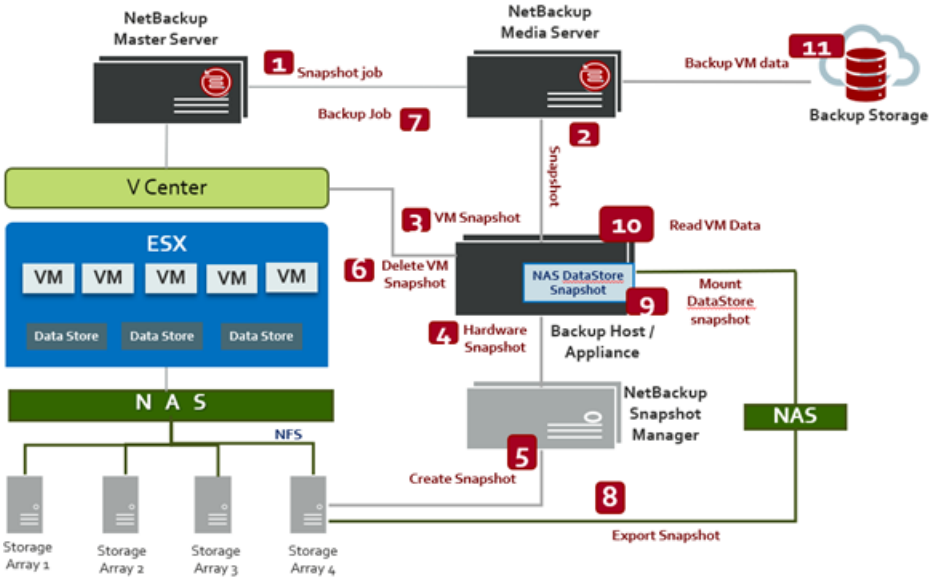

Note: This solution supports only the VMware datastores which are created on the NAS storage. It does not support the VMware datastores created on the SAN storage.

For all the supported NAS storage arrays, refer to the *NetBackup Snapshot Manager* section, under *Snapshot Solutions* in the *NetBackup Hardware and Cloud Storage Compatibility List (HCL)*.

Features and applications supported

Hardware snapshot-based protection for VMware includes the following features for protecting the virtual machine snapshots and replicated copies:

- Creates an instantaneous hardware snapshot of virtual machines.

- Backs up the virtual machines from the snapshots at primary locations and from replicated snapshots at remote locations.
- Block level incremental backup (BLIB) of the virtual machines from the snapshots at primary locations and from replicated snapshots at remote locations.
- Accelerator enabled backups of the virtual machines from the snapshots at primary locations and from replicated snapshots at remote locations.
- Supports browsing of virtual machine snapshots.
- Restores a virtual machine from its vmdk files that are in a snapshot.
- Restores an individual vmdk that is present in a snapshot.
- Restores the individual files from the vmdk files in a snapshot.
- Supports the storage lifecycle policies (SLPs).
- Under the Application Protection, following applications are supported in the VMware policy:
  - Microsoft Exchange databases
  - Microsoft SQL server

Prerequisites for hardware snapshot and replication

Following are the prerequisites for hardware snapshot-based support explained in the table.

Table 21-1 Prerequisites for hardware snapshot support

Support parameter	Description

System	■ All supported NetBackup primary, media server platforms.
	■ Backup host for VMware must be RHEL, SUSE or Windows.
	■ Snapshot manager server supported to the operating system platform as follows:
	  ■ Ubuntu 16.04 and 18.04 Server LTS
	  ■ Red Hat Enterprise Linux (RHEL) 8.2 and 7.x

Table 21-1 Prerequisites for hardware snapshot support *(continued)*

Support parameter	Description

Configuration	■ NetBackup version 10.1 primary, media server and backup host.
	■ VMware backup host can be on any of the NetBackup appliance form factor:
	  ■ NBA
	  ■ Flex
	  ■ NetBackup FlexScale (NBFS)
	■ NetBackup Snapshot Manager version 10.1

Permission	■ On Windows backup host, the following NetBackup services must be started using similar domain user account.
	  ■ NetBackup Client Service
	  ■ NetBackup Legacy Network Service
	■ The domain user must be part of local administrative group.

VMware NFS datastores	VMware NFS datastores mounted on the ESX host must be version NFS 4.1. or NFS 3.0.
VMware VCenter and ESX sever hosting virtual machine	Virtual machines must reside on the NFS datastores.

Operations supported with hardware snapshot

Table 21-2 Virtual machine operations with hardware snapshot

Operation	Description and notes
Create array-based snapshots of virtual machines on the NFS datastore.	Configure a storage lifecycle policy (SLP) and a backup policy to create array snapshots of virtual machines. The snapshots remain on the array or filer and are not backed up to a NetBackup media server storage unit. **Note**: ■ The snapshots are created on the NFS datastores only. ■ The virtual machine or its individual files can be restored directly from the snapshots on the storage array. The snapshots can also be replicated to other locations. ■ For fast browsing of files to restore, include the Index From Snapshot option in the SLP. This option, catalogs the metadata of the virtual machine.
Back up quiesce virtual machines from a snapshot (or snapshot replica) which resides on the NFS datastore.	Configure SLP and backup policy to make a backup image from the virtual machine snapshot. NetBackup backs up only the virtual machines quiesce before the snapshot occurs. The backup image is written to NetBackup storage unit. The image is retained according to the policy's retention period. **Note:** The **Application consistent snapshot** option in the policy must be enabled (Under **Options** > **Snapshot Client Options**).
Restore a virtual machine from a snapshot (or snapshot replica) that is on the NFS datastore or from the backup image written to NetBackup storage unit.	Use the NetBackup web UI interface to restore the virtual machine. Supported restore destinations are the original (NFS) datastore or an alternate datastore (NFS or non-NFS).

Table 21-2 Virtual machine operations with hardware snapshot *(continued)*

Operation	Description and notes
Restore individual files and VMDK from a snapshot (or snapshot replica) that is on the NFS datastore or from the backup image written to NetBackup storage unit.	Use the **Backup**, **Archive**, and **Restore** interface to restore the files. **Note**: ■ To restore files from a replica of the snapshot, the replica must exist in the same NetBackup domain as the snapshot. ■ To restore files to the original virtual machine, a NetBackup client must be installed on the original virtual machine.
Index from Snapshot	The Index From Snapshot operation catalogs the metadata of the virtual machine. This allows fast browsing of files to restore. Use Index From Snapshot option in the SLP to use this feature.
Live Browse of snapshot	The live-browse function allows you to browse the content of VM snapshot that resides on the storage array.

Configuring a VMware policy to use hardware snapshots

The following procedure describes how to configure a VMware policy using the NetBackup web UI to create hardware snapshots of the virtual machines that reside on an NFS datastore.

Note: This feature is not supported with protection plans for VMware. You need to create a VMware policy using the NetBackup web UI to use this feature.

For more information about configuring VMware policies, see the following:

See “Configure a VMware policy” on page 87.

See “About automatic virtual machine selection for NetBackup for VMware” on page 115.

Table 21-3 Configuration steps and description

Step	Description	Reference
1	Configure the NetBackup Snapshot Manager server in NetBackup.	*NetBackup Snapshot Manager for Data Center Administrator's Guide*
2	Configure the NAS storage array plug-in.	*NetBackup Snapshot Manager for Data Center Administrator's Guide*
3	Add the VMware backup host to your NetBackup configuration.	See "Add a VMware access host" on page 80.
4	Configure NetBackup access credentials for the VMware vCenter Server or ESX server.	See "Add VMware servers" on page 66.
5	Configure the SLP to use snapshot.	For more details, refer to these chapters in the *NetBackup Snapshot Manager for Data Center Administrator's Guide*,topic: *Configuring storage lifecycle policies for snapshots and snapshot replication*.
6	Configure a NetBackup VMware policy to perform the operations that are specified in the SLP.	

Only those policy options that are necessary to configure VMware policy to use hardware snapshot of VMs residing on NFS Datastore are listed in the following procedure.

To create a policy to use VM hardware snapshot

1. Open the NetBackup web UI.
2. On the left, select **Protection > Policies**.
3. Select **Add**.
4. Configure the options on the policy **Attributes** tab.
   - Enter a **Policy name** and from the **Policy type** list select **VMware**.
   - **Policy storage**: Select the SLP which you want to use and is configured for the snapshot-based protection.
   - **Perform snapshot backups**: Enable this option to automatically select other options which are required for the snapshot backup.
   - **Use Accelerator**: Select this option to accelerate backup operations.

Note: To accelerate backups, Backup From Snapshot must be defined in the SLP.

The MSDP storage unit that is used for the **Backup From Snapshot** operation must be the same as the MSDP storage unit that is used in the Snapshot (or snapshot replication).

5 Select **Snapshot options**.

- **Perform snapshot backup options**: Select **Snapshot options** to view the configuration parameters:
  - **Snapshot Type**: Select the appropriate snapshot type. By default, the **Auto** option is selected which enables NetBackup to automatically determine the snapshot type to be used for an array snapshot.
  - **Snapshot Manager**: Select the *NetBackup Snapshot Manager* host which communicates with the storage array to perform the snapshot operations.

    Note: To view the list of configured NetBackup Snapshot Manager hosts, users must have the **View** permission for Snapshot Managers in the RBAC role (**Global permissions > NetBackup Management > Snapshot Managers > View**).

  - **Maximum Snapshots**: Sets the maximum number of Snapshots to be retained at one time. When the maximum is reached, snapshot rotation occurs:
    The next snapshot causes the oldest to be deleted. Managed by SLP retention is automatically selected if the Fixed or the Expire after Copy retention is currently selected in the SLP.
  - **Application Consistent Snapshot**: This option is enabled by default. In most cases, NetBackup recommends that you keep this option enabled. To allow the SLP to create a backup image from the snapshot, this option must be enabled.
    If this option is *disabled*, data in the virtual machine may not be in a consistent state when the snapshot occurs. The snapshot may not capture all the data in the virtual machine. Also note that the following settings are disabled on the **VMware** tab: **Enable block-level incremental backup**, **Exclude deleted blocks**, **Exclude swap and paging files**, and **Application Protection**.

6 Select the **Schedules** tab. To configure the full and the incremental schedule for the backup, select **Add**.

Note: To use an incremental schedule, the **Enable block-level incremental backup** option must be enabled.

---

7 Use the **Clients** tab to create a query for the automatic selection of virtual machines. The selected VMs must reside on the NFS datastore.

See "Configure automatic virtual machine selection" on page 125.

8 Use **VMware** tab to select the virtual machine backup options.

Select **Enable block-level incremental backup** option to perform block level incremental backups.

Note: The **Transport modes** are not supported and are disabled. NetBackup uses the VMware file transport mode to move the data between the backup host and the storage array.

Note: Under the **Application protection** options only Microsoft Exchange and Microsoft SQL Server are supported.

---

9 When the policy configuration is complete, select **Create**.

Configuring a VMware policy to use NetBackup snapshot manager replication

Using the NetBackup Snapshot Manager for Data Center you can replicate the hardware snapshots of VMs. The replicated snapshots are accessed on VMware backup hosts to create point in time backup copies of VMs. The following procedure describes how to configure a VMware policy to use hardware snapshots and replication of VMs residing on NFS datastore.

See "Configure a VMware policy" on page 87.

See "About automatic virtual machine selection for NetBackup for VMware" on page 115.

Table 21-4 Configuration steps with description, and reference topics

Step	Description	Reference topic
1	Register the Snapshot Manager server in NetBackup.	For more details, see the *NetBackup Snapshot Manager for Data Center Administrator's Guide*.
2	Configure the NAS storage array plug-in.	For more details, see the *NetBackup Snapshot Manager for Data Center Administrator's Guide*.
3	Add the VMware access host to your NetBackup configuration.	See "Add a VMware access host" on page 80.
4	Configure NetBackup access credentials for the VMware vCenter Server or ESX server.	See "Add VMware servers" on page 66. See "Validate and update VMware server credentials" on page 71.
5	Configure the SLP to use snapshot and replication.	For more details, refer to these chapters in the *NetBackup Snapshot Manager for Data Center Administrator's Guide:* ■ Storage array replication ■ Configuring storage lifecycle policies for snapshots and snapshot replication ■ Supported storage array in datacenter
6	Configure a NetBackup VMware policy to perform the operations that are specified in the SLP.	See the procedure *To create a policy to use VM hardware snapshot in the web UI* in the following topic: See "Configuring a VMware policy to use hardware snapshots" on page 294.

Jobs in the Activity Monitor that use hardware snapshot for VMs

You can use the NetBackup Activity Monitor to keep track of virtual machines backups as they occur. The number of jobs that appear in the Activity Monitor depends on the policy's **Application Consistent Snapshot** option.

Note: By default, the **Application Consistent Snapshot** option is enabled. In most cases, NetBackup recommends that you keep this option enabled. If this option is disabled, data in the virtual machine may not be in a consistent state when the snapshot occurs.

Table 21-5 Job flow in the Activity Monitor

Application consistent snapshot option	Job flow in the Activity Monitor
Enabled	The first job discovers the virtual machines. This job is labeled Backup. Backup job starts with the following: ■ A Snapshot job for each virtual machine. ■ A Snapshot job for each datastore.
Disabled	The first job discovers the virtual machines. This job is labeled Backup. Backup job starts with the following: ■ A Snapshot job to collect all the virtual machines' configuration data. ■ A Snapshot job for each datastore.

Example 1: Virtual machine jobs with the **Application Consistent Snapshot** option enabled.

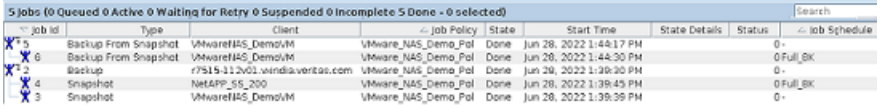

The jobs occurred as follows:

- The discovery (parent) Backup job for virtual machine discovery is ID 2.
- Job 3 made VMware snapshots of the virtual machine VMwareNAS_DemoVM.
- Job 4 made snapshots of datastore NetAPP_SS_200.
- Job 5 parent Backup from Snapshot export and mount the snapshot.
- Job 6 child Backup from Snapshot does the backup and creates the backup image.

Example 2: Virtual machine jobs with the **Application Consistent Snapshot** option disable.

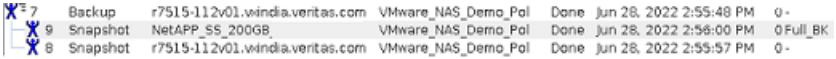

The jobs occurred as follows:

- The discovery (parent) Backup job for virtual machine discovery is ID 7.
- Job 8 collected the configuration data of all the virtual machines selected (VM1, VM2, and so forth).
- Jobs 9 snapshots of the virtual machine datastores.

Notes and limitations

Note the following about VMware NAS hardware snapshot for virtual machines datastores:

- Live browse from snapshot for XFS file system is not supported
- Single File Restore (SFR) from snapshot for XFS filesystem is only supported when Index from Snapshot step is present in the SLP.
- Agentless Single File Restore (ALVR) is not supported from the NetBackup web UI.
- GRT and Individual VMDK restore is not supported from the NetBackup web UI.
- While restoring the full VM from the incremental snapshot copy, the restore is performed only from the snapshot which is taken during the incremental backup. In the case of restore from the incremental backup image copy, the restore is performed from all the incremental images and the full backup image.
- To restore from the incremental backup images, all the primary copies must be either snapshot copies or the backup images copies.
- VMware policy which protects Microsoft Exchange using hardware snapshot-based backups, in that policy only Windows must be specified as the backup host.

Troubleshooting with VMware hardware snapshot and replication operations

About gathering information and checking logs

To create detailed log information, place a VERBOSE entry in the `bp.conf` file on the NetBackup primary and client. Or set the global logging level to a high value in the **Logging** properties for both the primary server and the client.

These directories can eventually require a lot of disk space. Delete them when you are finished troubleshooting and remove the VERBOSE option from the `bp.conf` file. Or reset the **Global logging level** to a lower value.

Logging directories for Linux platform

To create logging directories use the script `/usr/openv/netbackup/logs/mklogdir`. You can also create the directories using an access mode of `755` so NetBackup can write to the logs.

Table 21-6 Linux logging directories for snapshot operation

Path of log directory	Where to create the directory
/usr/openv/netbackup/logs/bprd	NetBackup primary server
/usr/openv/logs/nbjm	NetBackup primary server
/usr/openv/netbackup/logs/bpbrm	NetBackup media server
/usr/openv/netbackup/logs/bpfis	Backup host client

Table 21-7 Linux logging directories for backup operation

Path of log directory	Where to create the directory
/usr/openv/netbackup/logs/bprd	NetBackup primary server
/usr/openv/logs/nbjm	NetBackup primary server
/usr/openv/netbackup/logs/bpdbm	NetBackup primary server
/usr/openv/netbackup/logs/bptm	NetBackup media server
/usr/openv/netbackup/logs/bpbrm	NetBackup media server
/usr/openv/netbackup/logs/bpfis	Backup host client
/usr/openv/netbackup/logs/bppfi	Backup host client

Table 21-7 Linux logging directories for backup operation *(continued)*

Path of log directory	Where to create the directory
/usr/openv/netbackup/logs/bpbkar	Backup host client
/usr/openv/netbackup/logs/bppfi	Backup host client
/usr/openv/netbackup/logs/vxms	Backup host client

Table 21-8 Linux logging directories for single file restore operation

Path of log directory	Where to create the directory
/usr/openv/netbackup/logs/bprd	NetBackup primary server
/usr/openv/netbackup/logs/bpcd	Restore host client
/usr/openv/netbackup/logs/bpbkar	Restore host client
/usr/openv/netbackup/logs/bpfis	Restore host client
/usr/openv/netbackup/logs/bppfi	Restore host client
/usr/openv/logs/ncfnbhfr	Restore host client
/usr/openv/netbackup/logs/vxms	Restore host client
/usr/openv/netbackup/logs/tar	Destination client where the files are restored.

Table 21-9 Linux logging directories for full VM restore operation

Path of log directory	Where to create the directory
/usr/openv/netbackup/logs/bprd	NetBackup primary server
/usr/openv/netbackup/logs/bpbrm	NetBackup media server
/usr/openv/logs/ncfnbvmcopyback	Restore host client
/usr/openv/netbackup/logs/vxms	Restore host client

Table 21-10 Linux logging directories for live browse operation

Path of log directory	Where to create the directory
/usr/openv/netbackup/logs/bprd	NetBackup primary server
/usr/openv/netbackup/logs/bpdbm	NetBackup primary server
/usr/openv/netbackup/logs/bpcd	Backup host client

Table 21-10 Linux logging directories for live browse operation *(continued)*

Path of log directory	Where to create the directory
/usr/openv/netbackup/logs/bppfi	Backup host client
/usr/openv/logs/ncfnbbrowse	Backup host client
/usr/openv/netbackup/logs/vxms	Backup host client

Table 21-11 Linux logging directories for index from snapshot operation

Path of log directory	Where to create the directory
/usr/openv/logs/ncflbc	Backup host client
/usr/openv/netbackup/logs/vxms	Backup host client

Logging folders for Windows platforms

Table 21-12 Windows logging directories for snapshot operation

Path of log directory	Where to create the directory
/usr/openv/netbackup/logs/bprd	NetBackup primary server
/usr/openv/netbackup/logs/bpdbm	NetBackup primary server
/usr/openv/netbackup/logs/bpcd	NetBackup media server
/usr/openv/netbackup/logs/bppfi	Backup host client

Table 21-13 Windows logging directories for backup operation

Path of log directory	Where to create the directory
install_path\NetBackup\logs\bprd	NetBackup primary server
install_path\NetBackup\logs\nbjm	NetBackup primary server
install_path\NetBackup\logs\bpdbm	NetBackup primary server
install_path\NetBackup\logs\bptm	NetBackup media server
install_path\NetBackup\logs\bpbrm	NetBackup media server
install_path\NetBackup\logs\bpfis	Backup host client
install_path\NetBackup\logs\bppfi	Backup host client
install_path\NetBackup\logs\bpbkar	Backup host client

Table 21-13 Windows logging directories for backup operation *(continued)*

Path of log directory	Where to create the directory
install_path\NetBackup\logs\bppfi	Backup host client
install_path\NetBackup\logs\vxms	Backup host client

Table 21-14 Windows logging directories for single file restore operation

Path of log directory	Where to create the directory
install_path\NetBackup\logs\bprd	NetBackup primary server
install_path\NetBackup\logs\bpcd	Restore host client
install_path\NetBackup\logs\bpbkar	Restore host client
install_path\NetBackup\logs\bpfis	Restore host client
install_path\NetBackup\logs\bppfi	Restore host client
install_path\NetBackup\logs\ncfnbhfr	Restore host client
install_path\NetBackup\logs\vxms	Restore host client
install_path\NetBackup\logs\tar	Destination client where the files are restored.

Table 21-15 Windows logging directories for full VM restore operation

Path of log directory	Where to create the directory
install_path\NetBackup\logs\bprd	NetBackup primary server
install_path\NetBackup\logs\bpbrm	NetBackup media server
install_path\NetBackup\logs\ncfnbvmcopyback	Restore host client
install_path\NetBackup\logs\vxms	Restore host client

Table 21-16 Windows logging directories for live browse operation

Path of log directory	Where to create the directory
install_path\NetBackup\logs\bprd	NetBackup primary server
install_path\NetBackup\logs\bpdbm	NetBackup primary server
install_path\NetBackup\logs\bpcd	Backup host client
install_path\NetBackup\logs\bppfi	Backup host client

Table 21-16 Windows logging directories for live browse operation *(continued)*

Path of log directory	Where to create the directory
install_path\NetBackup\logs\ncfnbbrowse	Backup host client
install_path\NetBackup\logs\vxms	Backup host client

Table 21-17 Windows logging directories for index from snapshot operation

Path of log directory	Where to create the directory
install_path\NetBackup\logs\ncflbc	Backup host client
install_path\NetBackup\logs\vxms	Backup host client

Chapter 22

Best practices and more information

This chapter includes the following topics:

NetBackup for VMware best practices

The following are best practices for NetBackup for VMware:

- For a more efficient backup, the NetBackup media server and the VMware backup host should be installed on the same host.
- When creating virtual machines, use the same name for both host name and display name. If the policy's **Primary VM identifier** option is changed, the existing entries on the policy **Clients** tab still work.
- VMware recommends that you run no more than four simultaneous backups of virtual machines that reside on the same datastore.
- Successful VMware snapshots depend on the following:
  - The amount of I/O that occurs on the virtual machine datastore. Backups should be scheduled when relatively little I/O activity is expected. Reducing the number of simultaneous backups can also help.

Limit access to the datastore per policy: Use the **Limit jobs per policy** attribute in the NetBackup policy.

Limit access to the datastore globally (across all policies): Use the Host Properties **Resource Limit** screen.

See "Change resource limits for VMware resource types" on page 81.

- The design of the I/O substructure that is associated with each virtual machine datastore. For correct I/O design and implementation, consult your VMware documentation.

- Make sure that the VMware backup host has enough memory to handle the number of simultaneous backups that occur.
- Include in a single NetBackup policy those virtual machines that use the same datastore. This practice lets you control the amount of backup-related I/O that occurs per datastore, to limit the backup effect on the target virtual machines.
- NetBackup supports multiple backup hosts. When a single backup host is saturated with a backup process, another backup host can be added to increase backup throughput.
- If a VM's disks are accessible to multiple ESX hosts, the disks can be accessed for backup or restore through any of the ESX hosts. The ESX host may or may not be the ESX host where the virtual machine is running or registered. All of the following must be accessible to each other and should have DNS configured:
  - The vCenter server.
  - All ESX hosts under the vCenter that have access to the VM's vmdk files.
  - The backup host.
- Upgrade to the latest version of VMware vSphere or Virtual Infrastructure.

NetBackup for VMware with deduplication

For a VMware backup to a deduplication storage unit, select the **Enable file recovery from VM backup** option on the **VMware** policy tab. This option provides the best deduplication rates. Without the **Enable file recovery from VM backup** option, the result is a lower rate of deduplication.

More information is available on the VMware options.

See "Backup options on the VMware tab" on page 90.

How NetBackup handles VMware tag associations at restore

To better understand the NetBackup restore behavior, it's important to understand some of the internal mechanics of how VMware handles tag associations. For each tag that is created in VMware, there is a corresponding internal identifier you cannot see or edit. This feature allows VMware to function correctly without the need to account for variation in naming conventions. Tags are replicated throughout the environment through the Platform Services Controller (PSC).

NetBackup recognizes and uses tag names when they are part of the VIP query. Backups of the virtual machine store the tag name and VMware internal identifier for all tags that are associated with that virtual machine. During the restore, however, NetBackup only creates the tag associations based on the VMware internal identifiers that are defined on the target vCenter server.

Figure 22-1 Sample VMware environment

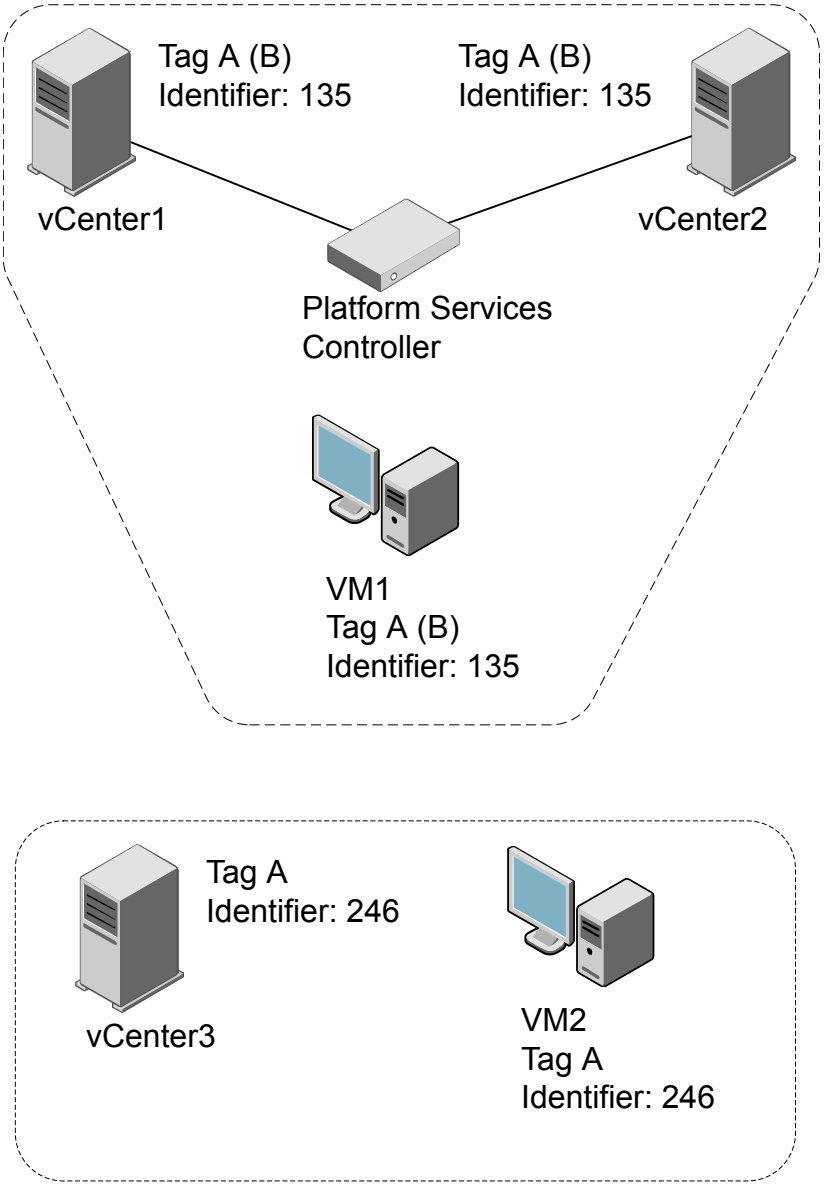

Figure 22-1 shows a sample VMware environment with multiple vCenter servers, virtual machines, and tags. Assume that a backup of all virtual machines completes successfully with a NetBackup status code 0.

- If you restore VM1 to either vCenter1 or vCenter2, it is restored with tag **A** and the restore exits with a NetBackup status code 0. This behavior is true both for a restore to the same name as well as an alternate client restore. This behavior is the result of VMware replicating tags across all vCenter servers that are attached to a single PSC.
- If you restore VM1 to vCenter3, it is restored without any tags. The restore exits with a NetBackup status code 1. This behavior is because VMware uses the internal identifier. While there is a tag name **A** in vCenter3, the internal identifier for tag **A** does not match the internal identifier that is restored. This behavior is true both for a restore to the same name as well as an alternate client restore.
- After the NetBackup backup, if tag **A** is renamed to **B**, when VM1 is restored to either vCenter1 or vCenter2 it is restored with tag **B**. The restore exits with a NetBackup status code 0. This behavior is because VMware uses the internal identifier, and now associates this identifier with the tag name **B**.
- After the NetBackup backup, if tag **A** is deleted, when VM1 is restored to either vCenter1 or vCenter2 it is restored without any tag associations. The restore exits with a NetBackup status code 1.

If for any reason the backup of the virtual machines in Figure 22-1 did not successfully capture the tag associations, the backup exits with NetBackup status code 0. The reasons for failing to capture tag associations appear in the Activity Monitor. Any restores based on this backup exit with NetBackup status code 0, but no tag information is restored. Depending on the backup error, more information regarding tag associations may appear in the Activity Monitor.

Best practices for VMware tag usage

Use unique tag names whenever possible

The combination of **Category** and **Tag** uniquely identifies a tag association to a virtual machine. Cohesity, however, does not support selection of virtual machines based on **Category**. As such, avoid the creation of tags with identical names in different **Categories**.

When you create a VMware Intelligent Policy and select virtual machines based on tags, identical tag names in different categories may have unintended consequences. All virtual machines with the tag name are selected independent of the **Category**. Be aware of this behavior as you create your VMware Intelligent Policies.

Use a primary server or a media server as a discovery host for any VMware intelligent policies that use tags

Because tag queries require Java, the discovery host that is used for VMware intelligent policies must have Java installed. NetBackup primary servers have Java installed by default. Java is an optional component for NetBackup media servers and for UNIX and Linux clients, and may or may not be installed in your environment. Windows clients that are used as a discovery host require a separate installation of the Remote Console.

Successful use of tags in a mixed vCenter Server environment

NetBackup support for VMware tags begins at vCenter 6.0. In a mixed vCenter Server environment (for example, 5.x and 6.0), a VMware intelligent policy query that uses the **Tag** keyword can return some virtual machines as **Failed** if the configuration of the query requires tag evaluation on a 5.x vCenter. This behavior is observed when you select **Test query** during policy configuration. This behavior is also observed when you run the policy. The parent (discovery) job exits with a non-zero status and its details enumerate the virtual machines that **Failed**. The virtual machines that the query has included are backed up normally.

You can use any of the options that are shown to create a query that uses tags but does not report virtual machines on 5.x servers as **Failed**.

- Use another field to limit tag evaluation to the supported vCenter Server versions.
  Example: `vCenterVersion GreaterEqual "6.0.0" AND Tag Equal "Production"`
  Example: `vCenter Equal "vcenter-123" AND tag Equal "Production"`
- Use the **VMware server list** option under the **VMware advanced attributes** on the **VMware** tab to restrict the policy to a supported list of vCenter Servers.
  See Table 7-6 on page 97.
- Enable the option **Treat tags as unset if unable to evaluate** under the **VMware advanced attributes** on the **VMware** tab.
  See Table 7-6 on page 97.

Place a tag clause at the end of the VMware intelligent policy query

Place the tag clause toward the end of the VMware intelligent policy query. This configuration allows NetBackup to eliminate as many virtual machines as possible before it evaluates the tag portion of the query. This organization optimizes the performance of discovery and selection of virtual machines.

Consider the query: `Powerstate Equal poweredOn AND Tag Equal "Production"`

In the example, the first clause automatically eliminates all virtual machines that are `poweredOff`. The query does not need to evaluate the tag clause of the query for all of those virtual machines.

Tag backup and restore are a best effort

Tag associations are part of the metadata of the virtual machine. NetBackup considers virtual machine tag association protection a best effort backup. Any tag collection errors are shown in the Activity monitor for the virtual machine snapshot job.

About reducing the size of VMware backups

NetBackup provides the following options for reducing the backup size for a VMware virtual machine:

Table 22-1 Options for reducing the virtual machine backup size

Option	Description
Block level incremental backup (BLIB)	BLIB reduces the size of backups (full and incremental) by tracking block-level changes. Only the blocks that have changed since the last full or incremental are included in the backup. For incremental backups, this option applies to cumulative and to differential backups. BLIB works with VMware's Changed Block Tracking in vSphere to track block-level changes in the virtual machine. The **Enable block-level incremental backup** option is enabled by default on the NetBackup policy **VMware** tab. NetBackup uses BLIB for storage optimization when the backup runs. Note the following: ■ Storage optimization cannot be used if a snapshot exists on the virtual machine when VMware Changed Block Tracking is turned on. ■ The first backup you run with BLIB must be a full backup. See "Block-level backup (BLIB): full vs incremental" on page 313. ■ BLIB works only with ESX 4.0 or later virtual machines at version vmx-07 or later. ■ If you used vSphere Client to manually create a VM snapshot and that snapshot is currently active, you may have to delete the snapshot. See "Deleting a vSphere Client snapshot" on page 313.
Exclusion of deleted blocks.	Reduces the size of virtual machine backups by excluding any deleted sectors in the file system on the virtual machine. To enable this option, click **Exclude deleted blocks** on the policy **VMware** tab. Refer to **Exclude deleted blocks** in the following topic: See "Optimizations options (VMware)" on page 91.

Block-level backup (BLIB): full vs incremental

When you use the **Enable block-level incremental backup** option in the policy, NetBackup uses VMware's Changed Block Tracking feature (CBT) to reduce the backup size.

This option reduces the size of full backups as well as the size of incremental backups, as follows.

Table 22-2 Block-level backup of the virtual machine: full vs incremental schedule

Type of backup	Optimization that is used in backup
Backup of entire virtual machine, with full schedule	Backs up only the blocks that have changed since the .vmdk was created. Note that the blocks that are not initialized are excluded from the backup.
Backup of entire virtual machine, with incremental schedule	Backs up only the blocks that have changed since the last backup, as follows: ■ For cumulative incrementals, BLIB backs up only the blocks that changed since the last full backup. ■ For differential incrementals, BLIB backs up only the blocks that changed since the previous backup of any kind. **Note:** On incremental backups, the **Enable block-level incremental backup** option backs up the changed files as well as their metadata. **Note:** VMware CBT may occasionally reset tracking of file changes in the virtual machine, such as after a power failure or hard shutdown. For the next backup, NetBackup reads all the data from the vmdk files and the backup takes longer than expected. If deduplication is enabled, the deduplication rate is lower than expected. The following VMware article contains more information on CBT: Changed Block Tracking (CBT) on virtual machines (1020128)

Deleting a vSphere Client snapshot

To use BLIB with NetBackup for VMware, you must delete an existing vSphere Client snapshot if both of the following are true:

- You used the vSphere Client interface to manually create a snapshot of the virtual machine and that snapshot is currently active.

- A NetBackup policy with the **Enable block-level incremental backup** feature had never been used to back up the virtual machine before you started the vSphere Client snapshot.

In this case, NetBackup cannot enable BLIB for the virtual machine. You must delete the vSphere Client snapshot. Then, when the NetBackup policy runs, BLIB is enabled.

More information is available on NetBackup block-level backups:

To delete a vSphere Client snapshot

1. In the vSphere Client interface, right-click on the virtual machine and select **Snapshot > Snapshot Manager**.
2. Select the snapshot and click **Delete**.
3. To back up the virtual machine with BLIB, you can now run a NetBackup policy that uses **Enable block-level incremental backup**.

   NetBackup backups from this policy continue to use storage optimization, even if you manually create a snapshot of the virtual machine using vSphere Client.

Further assistance with NetBackup for VMware

Table 22-3 Sources of information on NetBackup for VMware

Topic	Source
Snapshot Client configuration	NetBackup Snapshot Manager for Data Center Administrator's Guide.
List of supported combinations of platforms and snapshot methods	See the Snapshot Client section in the Hardware and Cloud Storage Compatibility List (HCL)
Support information on NetBackup for VMware	NetBackup Software Compatibility List

Chapter 23

Troubleshooting VMware operations

This chapter includes the following topics:

- Conflict between NetBackup and VMware Storage vMotion with vSphere 5.0 or later
- Backup or restore job hangs
- VMware SCSI requirements for application quiesce on Windows
- VMware virtual machine does not restart after restore
- A restored VM may not start or its file system(s) may not be accessible
- NetBackup job fails due to update tasks on the VMware server
- The vSphere interface reports that virtual machine consolidation is needed
- Linux VMs and persistent device naming
- For a VMware virtual machine with Windows dynamic disks, a restore from incremental backup fails with a Windows restore host and the hotadd transport mode
- Simultaneous hotadd backups (from the same VMware backup host) fail with status 13
- Troubleshooting VMware tag usage
- Ensuring that guest customizations can be restored in vCloud Director
- Troubleshooting vmdk restore to existing VM
- Troubleshooting backups of virtual machines on Virtual Volumes (VVols)
- Issues with the CA certificate during installation of the NetBackup client on VMware Cloud (VMC)

NetBackup logging for VMware

For log messages about VMware backup or VMware restore, see the following NetBackup log directories.

Table 23-1 NetBackup logs that pertain to VMware backup and restore

Log directory	Contains the messages on	Resides on
install_path\NetBackup\logs\bpbrm	Backup and restore	NetBackup primary or media server
install_path\NetBackup\logs\bptm	Backup and restore	NetBackup media server

Table 23-1 NetBackup logs that pertain to VMware backup and restore *(continued)*

Log directory	Contains the messages on	Resides on
install_path\NetBackup\logs\bpfis	Snapshot creation and backup	VMware backup host
install_path\NetBackup\logs\bpcd	Snapshot creation and backup	VMware backup host
install_path\NetBackup\logs\bpbkar	Backup	VMware backup host
install_path\NetBackup\logs\bprd	Restore	NetBackup primary server
install_path\NetBackup\logs\nbfsd	Instant recovery	NetBackup media server and VMware backup host
install_path\NetBackup\logs\bpVMutil	Policy configuration and on restore	VMware backup or recovery host
install_path\NetBackup\logs\bpVMreq	Restore	The client where the Backup, Archive, and Restore interface is running.
install_path\NetBackup\logs\nbproxy	Policy configuration	VMware backup host
install_path\NetBackup\logs\ncfnbcs (originator ID 366) ncfnbcs uses unified logging. See the *NetBackup Logging Reference Guide* for information on how to use unified logs.	Automatic virtual machine selection, and disabling and re-enabling VMware Storage vMotion during backup or restore.	VMware backup host
install_path\NetBackup\logs\ncfnbrestore (originator ID 357) ncfnbrestore uses unified logging. See the *NetBackup Logging Reference Guide* for information on how to use unified logs.	Restore	VMware recovery host
Windows: *install_path*\NetBackup\logs\vxms Linux: /usr/openv/netbackup/logs/vxms	File mapping during backup	VMware backup host See "Configuring VxMS logging" on page 319. **Note:** The use of VxMS logging can reduce the performance of the backup host.

Note: Except for unified logging directories, these log directories must already exist in order for logging to occur. If these directories do not exist, create them.

To create most of these log directories, run the following command on the NetBackup servers and backup host:

Windows:

`install_path\NetBackup\logs\mklogdir.bat`

UNIX (on primary or media servers):

`/opt/openv/netbackup/logs/mklogdir`

See "Configuring VxMS logging" on page 319.

More detail is available on snapshot logs and logging levels. See the *NetBackup Snapshot Client Administrator's Guide*.

A broader discussion of NetBackup logging is available. See the NetBackup Snapshot Manager for Data Center Administrator's Guide.

NetBackup logs for Accelerator with virtual machines

Accelerator does not require its own log directory. For log messages about Accelerator, see the following standard NetBackup log directories.

Table 23-2 NetBackup logs that may contain Accelerator information

Log directory	Resides on
UNIX: `/usr/openv/netbackup/logs/bpbrm` Windows: `install_path\NetBackup\logs\bpbrm`	NetBackup primary or media server
UNIX: `/usr/openv/netbackup/logs/bptm` Windows: `install_path\NetBackup\logs\bptm`	NetBackup media server
`/usr/openv/netbackup/logs/bpbkar` Windows: `install_path\NetBackup\logs\bpbkar`	Backup host
`/usr/openv/netbackup/logs/bpfis` Windows: `install_path\NetBackup\logs\bpfis`	Backup host
VxMS logs	See "NetBackup logging for VMware" on page 316.

To create the log directories, run the following command on the NetBackup servers and backup host:

On Windows:

install_path\NetBackup\logs\mklogdir.bat

On UNIX/Linux:

/usr/openv/netbackup/logs/mklogdir

Configuring VxMS logging

The following procedures describe how to configure VxMS logging for NetBackup.

Note: VxMS logging may require significant resources on the VMware backup host.

VxMS logging on a Linux backup host

To configure VxMS logging on a Linux backup host

1. Create the VxMS log directory:

   /usr/openv/netbackup/logs/vxms

   Note: For logging to occur, the VxMS directory must exist.

   Note: If you have run the NetBackup `mklogdir` command, the VxMS log directory already exists.

   See "NetBackup logging for VMware" on page 316.

2. Add the following to the `/usr/openv/netbackup/bp.conf` file:

   VXMS_VERBOSE=*<numeric value of 0 or greater>*

   See Table 23-3 for the available logging levels.

3. To change the log location, enter the following in the `bp.conf` file:

   vxmslogdir=*path to new log location*

   Note: If the VxMS log location is changed, the Logging Assistant does not collect the logs.

VxMS logging on a Windows backup host

To configure VxMS logging on a Windows backup host

1. Create the VxMS log directory:

   install_path\NetBackup\logs\vxms

   Note: For logging to occur, the VxMS folder must exist.

   Note: If you have run the NetBackup `mklogdir.bat` command, the VxMS log directory already exists.

   See "NetBackup logging for VMware" on page 316.

2. In the Windows registry, create the DWORD registry entry `VXMS_VERBOSE` in the following location:

   HKEY_LOCAL_MACHINE > SOFTWARE > Veritas > NetBackup > CurrentVersion > Config

3. To configure the logging level, set the numeric value of `VXMS_VERBOSE` to 0 or greater. Larger numbers result in more verbose logs.

   See Table 23-3 for the available logging levels.

4. To change the log location:

   - Open regedit and go to the following location:
     HKEY_LOCAL_MACHINE > SOFTWARE > Veritas > NetBackup > CurrentVersion
   - Create the registry entry `vxmslogdir` with a string value `(REG_SZ)`. For the string value, specify the full path to an existing folder.

   Note: You can use NTFS compression on VxMS log folders to compress the log size. The new logs are written in compressed form only.

   Note: If the VxMS log location is changed, the Logging Assistant does not collect the logs.

VxMS logging levels

Table 23-3 lists the VxMS logging levels.

Note: Logging levels higher than 5 cannot be set in the Logging Assistant.

Note: Logging levels higher than 5 should be used in very unusual cases only. At that level, the log files and metadata dumps may place significant demands on disk space and host performance.

Table 23-3 VxMS logging levels

Level	Description
0	No logging.
1	Error logging.
2	Level 1 + warning messages.
3	Level 2 + informative messages.
4	Same as level 3.
5	Highly verbose (includes level 1) + auxiliary evidence files (`.mmf, .dump, VDDK logs, .xml, .rvpmem`). You can set the logging level for the VDDK messages.
6	VIX (VMware virtual machine metadata) dumps only.
7	VHD (Hyper-V virtual machine metadata) dumps only.
> 7	Full verbose + level 5 + level 6 + level 7.

Format of the VxMS core.log and provider.log file names

For the log files `core.log` and `provider.log` created by default during VxMS logging, the NetBackup administrator's user name is inserted into the log file name.

Table 23-4 describes the format of the log file names.

Table 23-4 Format of VxMS core.log and provider.log file names

Platform	VxMS log-file-name format

Windows	`VxMS-`*`thread_id-user_name`*`.`*`mmddyy_tag`*`.log`
	For example:
	`VxMS-7456-ALL_ADMINS.070214_core.log`
	`VxMS-7456-ALL_ADMINS.070214_provider.log`

UNIX, Linux	`VxMS-`*`thread_id-user_name`*`.log.`*`mmddyy_tag`*
	For example:
	`VxMS-27658-root.log.081314_core`
	`VxMS-27658-root.log.081314_provider`

See "Configuring VxMS logging" on page 319.

Configuring the VDDK logging level

The following NetBackup processes capture VDDK log messages:

- `bpbkar, bpbkarv, nbrestore`

  These processes write VDDK messages in the VxMS logs if the VxMS logging level (VxMS_VERBOSE) is 5 or higher. By default, when VXMS_VERBOSE is 5 or higher, the VDDK messages are generated at the highest verbosity.

- `ncfnbcs, bpVMutil`

  These processes write VDDK messages in their own log directories if the NetBackup global logging level is 5. By default, when the NetBackup global logging level is 5, the VDDK messages in the ncfnbcs and bpVMutil logs are generated at minimum verbosity.

To change the logging level (verbosity) of the VDDK messages

1. Check the following on the backup host:
   - Make sure the VxMS log directory exists and that the VXMS_VERBOSE DWORD is set to 5 or higher.
     See "Configuring VxMS logging" on page 319.
   - Make sure the bpVMutil log directory exists and that the NetBackup global logging level is set to 5.
     Note: The ncfnbcs process uses unified logging: You do not have to manually create a log directory for ncfnbcs. More information on unified logging is available in the NetBackup Logging Reference Guide.
2. Enter the following on the backup host:
   - Windows:
     Create the DWORD registry entry VDDK_VERBOSE in the following location:
     HKEY_LOCAL_MACHINE > SOFTWARE > Veritas > NetBackup > CurrentVersion > Config
   - Linux:
     Add the following to the `/usr/openv/netbackup/bp.conf` file:

```
VDDK_VERBOSE=numeric value
```

3 Set the numeric value of VDDK_VERBOSE as follows:

0	Panic (failure messages only).
1	Level 0 + error logging.
2	Level 1 + warning messages.
3	Level 2 + audit messages.
4	Level 3 + informational messages.
5	Highly verbose; level 4 + additional details.
6	Most verbose; level 5 + debug messages.

Troubleshooting VMware backups

The following table describes the issues that may occur when you perform VMware backups.

Table 23-5 Errors with VMware backups

Issue	Explanation

Backup fails with `Status 13`.	NetBackup allows up to 31 snapshots per virtual machine. If the virtual machine has more than 31 snapshots, the backup may fail with status 13. Messages similar to the following appear in the NetBackup job details:
	
	`10/18/2012 4:56:59 PM - Critical bpbrm(pid=4604)from client Umesh_w2k3_hypervm33: FTL - vSphere_freeze: Unable to remove existing snapshot, too many existing snapshots (91).`
	`10/18/2012 4:56:59 PM - Critical bpbrm(pid=4604) from client Umesh_w2k3_hypervm33: FTL - VMware_freeze: VIXAPI freeze (VMware snapshot) failed with 26: SYM_VMC_REMOVE_SNAPSHOT_FAILED`
	
	As a reminder to consolidate or delete snapshots, the NetBackup detailed status provides the following message when the number of snapshots exceeds 15:
	
	`Umesh_w2k3_hypervm33: WRN - vSphere_freeze: VM has 16 existing snapshots. Snapshots may start failing if this number exceeds 32`
	
	Recommended action:
	
	■ Consolidate or delete the existing snapshots. Then rerun the backup.
	■ Use Replication Director for any backups that require more than 31 snapshots per virtual machine.

Backup fails with message: `FTL - vSphere_freeze: Unable to proceed with snapshot creation, too many existing delta files(50)`	If the number of snapshot delta files for a VM's vmdk exceeds 32, snapshot creation fails. A message similar to the following appears in the NetBackup detailed status:
	
	`Umesh_w2k3_hypervm33: FTL - vSphere_freeze: Unable to proceed with snapshot creation, too many existing delta files(50).`
	
	As a reminder to consolidate or delete snapshots, the NetBackup detailed status provides the following message when a vmdk's delta files exceed 16:
	
	`Umesh_23k3_hypervm33: WRN - vSphere_freeze: VM has 17 existing delta files for vmdk Umesh_23k3_hypervm33.vmdk. Snapshots may start failing if this number exceeds 31`
	
	Delta files can accumulate if the VM's snapshots are not deleted or consolidated. Consolidate or delete the existing snapshots, then rerun the backup.

Table 23-5 Errors with VMware backups *(continued)*

Issue	Explanation
For an independent disk, the backup succeeds but the backup image contains no data for the independent disk.	NetBackup for VMware cannot back up the data on an independent disk, because an independent disk cannot be captured with a snapshot. To back up the data on an independent disk, install a NetBackup client on the virtual machine. You can configure NetBackup to back up the virtual machine and any independent disks as if the client was installed on a physical host. You can restore the virtual machine and then restore the independent disk as a separate job.
NetBackup fails the backups of virtual machines that are empty.	A VM may be empty because all of its disks are empty or because disk exclusion excludes all disks.
For the virtual machines that vSphere 6.5 hosts, NetBackup fails the backup if a snapshot exists while NetBackup tries to enable VMware Change Block Tracking.	Because the VMware Change Block Tracking API behavior has changed beginning in vSphere 6.5, NetBackup fails the backup. NetBackup enables CBT on a VM if **Block Level Incremental Backups** is enabled in the backup policy and CBT is not enabled already on the VM.
The status log displays: `There is no complete backup image match, a regular full backup will be performed.`	If a policy is changed from manual selection to Intelligent policy (or vice versa), the next backup of the VM is a regular full backup, even if a backup already exists for that VM.
The virtual machine cannot be quiesced in preparation for the snapshot.	Make sure that the VMware Tools are installed and up to date on each virtual machine.

Troubleshooting the restore of VMware and restores of files

The following table describes the issues that may occur when you perform VMware restores.

Table 23-6 Errors with VMware restores and file restores

Issue	Explanation
Restore fails because the datastore did not have enough space for the `.vmdk` files.	This issue can occur when a virtual machine is configured on multiple datastores and a leftover snapshot that existed on the virtual machine when it was backed up. NetBackup tries to restore all `.vmdk` files to the snapshot datastore. Alternatively, you can restore the virtual machine to an alternate location.
File recovery from a VM backup is unsuccessful.	For a Linux virtual machine, if unsupported special characters are in the volume name, the **Enable file recovery from VM backup** option does not work. As a result, you cannot restore individual files from that volume. The following topic for supported characters. See "Optimizations options (VMware)" on page 91.
	If a VM is configured on a logical volume (LVM or LDM): Individual file recovery does not work if the volume disk set contains a mixture of disk types: ■ Some of the logical volume disks are regular virtual disks (normal VMDKs). ■ Some of the disks in the same volume are independent disks or are physical disks in raw device mapping mode (RDM). The backup job succeeds but files cannot be individually restored from the file systems that reside on the disk set (LVM or LDM). To be able to restore files individually, reconfigure the VM's logical volumes to reside on regular virtual disks (vmdk) only. Note that VMware does not make snapshots of independent disks or RDM disks.
An incremental backup does not back up files and the individual files cannot be restored from the incremental backup.	Any files that are moved or renamed or not backed up. However, when you restore the entire VM from a block-level incremental backup, note: the file metadata is updated and the moved or renamed files in the restored VM reflect the updated metadata.
The restore fails when you restore individual files to a virtual machine that has NetBackup client software.	When you restore individual files to a virtual machine that has a NetBackup client, make sure that a firewall does not interfere with the restore. If a firewall stops the restore, turn off the firewall and retry the restore.
Mount point missing on a restored Windows virtual machine.	A Windows virtual machine may fail to write its mount point configuration to disk (the mount point configuration remains in RAM). In that case, the mount point information cannot be backed up. When the virtual machine is restored, the data from the mounted volume is restored, but the mount point is absent from the restored virtual machine. Reassign the mount point on the restored virtual machine. To make sure the mount point is correctly configured, restart the virtual machine.

Table 23-6 Errors with VMware restores and file restores *(continued)*

Issue	Explanation
Recovery of individual files or folders is not available and requires “Switch to Instant Access”.	In some cases you may find you cannot access or recover certain files with an individual file restore from a VMware backup. However, it may be possible to recover these files with the “Switch to Instant Access” feature in the web UI. Some examples of these files include files from unsupported file systems (for example, btrfs or thin-provisioned LVM volumes) or unsupported file system features (for example, files with XFS reflinks or shared extents). Additionally, if certain mount points do not display in the browse tree or list view, you may have to click “Switch to Instant Access“ to view these mount points. See “Recover files and folders with VMware agentless restore” on page 245.
Mount points are not available when restoring files from a Linux virtual machine.	For Linux virtual machines, only the ext2, ext3, ext4, and XFS file systems are supported for individual file restore. If a partition is formatted with some other file system, the backup succeeds but NetBackup cannot map the file system addresses of the files. As a result, NetBackup cannot restore individual files from that partition. Only the files that were on ext2, ext3, ext4, or XFS partitions can be individually restored. **Note:** To restore individual files from their original mount points, the "/" (root) partition must be formatted as ext2, ext3, ext4, or XFS. If the "/" (root) partition is formatted with a different file system (such as ReiserFS), the mount points cannot be resolved. In that case, you can restore ext2, ext3, ext4, or XFS files from the `/dev` level (such as `/dev/sda1`). You cannot restore the files from their original mount point level.
Invalid client error when you restore files using the BAR interface that is installed on the virtual machine.	If the virtual machine was backed up by display name or UUID, and the display name is not the same as the host name, note: You cannot restore individual files by means of the Backup, Archive, and Restore (BAR) interface if the interface is installed on the virtual machine itself. The files can be restored if BAR is installed on the primary server or media server. In this case, BAR must not be installed on the virtual machine that you want to restore to. To restore files, the **Destination client for restores** field in the BAR interface must have a valid host name or IP address.
An attempt to restore a full virtual machine fails with the SAN transport type.	Recommended action: Try the NBD transport type instead.

Table 23-6 Errors with VMware restores and file restores *(continued)*

Issue	Explanation
Restoring a virtual machine with a transport mode of NBD or NBDSSL is slow.	The virtual machine had many small data extents due to heavy fragmentation. (A file system extent is a contiguous storage area defined by block offset and size.) Recommended action: Use the hotadd transport mode.
	The restore is from a block-level incremental backup and the changed blocks on the disk were heavily fragmented when the incremental backup occurred. Recommended action: Use the hotadd transport mode.
For the SAN transport mode, the job is slow.	This issue can occur when you restore to a vCenter Server. Recommended action: For greater speed, designate a VMware restore ESX server as the destination for the restore. See "Add VMware servers" on page 66.
	For other circumstances, see the following article: VMware Transport Modes: Best practices and troubleshooting

For the SAN transport mode and a restore host on Windows, the restore fail.	The datastore's LUN is offline. The detailed status log contains messages similar to the following:
	`5/22/2013 4:10:12 AM - Info tar32(pid=5832) done. status: 24: socket write failed`
	`5/22/2013 4:10:12 AM - Error bpbrm(pid=5792) client restore EXIT STATUS 24: socket write failed`
	Recommended action:
	■ Make sure the status of the SAN disk on the restore host is online (not offline). Disk status can be checked or changed using the Windows diskpart.exe utility or the Disk Management utility (diskmgmt.msc). When the disk status reads online, retry the restore.
	■ If multipathing is enabled, make sure all the paths are online.

Restores that use the hotadd or SAN transport modes do not include the VM's metadata changes in the restore.	The status log of the NetBackup job contains messages similar to the following:
	`07/25/2013 12:37:29 - Info tar (pid=16257) INF - Transport Type = hotadd`
	`07/25/2013 12:42:41 - Warning bpbrm (pid=20895) from client <client_address>: WRN - Cannot set metadata (key:geometry. biosSectors, value:62) when using san or hotadd transport.`
	Recommended action: Retry the restore with a different transport mode (**nbd** or **nbdssl**).
	This problem is a known VMware issue.

Table 23-6 Errors with VMware restores and file restores *(continued)*

Issue	Explanation
You cannot restore individual VMware files onto the virtual machine itself, except under certain conditions.	Make sure that the VMware Tools are installed and up to date on each virtual machine.

Troubleshooting the adding of VMware servers

Table 23-7 Errors adding VMware servers

Issue	Explanation
Virtualization server credential validation fails.	This error occurs when the NetBackup primary server is in a DNAT or a similar setup can access only a few specified NetBackup hosts (`PROXY_SERVERS`). The credentials validation occurs in the following order: ■ The auto-discovered discovery host is used to access the virtualization server. ■ If the autodiscovery does not find any information about the virtualization server on the discovery host, the NetBackup primary server is used. Workaround: When you add the virtualization server credentials, select the proxy server that has access to the virtualization server as the backup host for validation. **Note:** Adding or updating VMware credentials also automatically starts the discovery of the VMware server. When backup host information is provided in the request, it is used to perform validation of credentials as well as for performing the discovery. For discovery, NetBackup 8.1.2 is the minimum version that is supported for a NetBackup media server or client that serves as a backup host. For older versions, backup host credential validation succeeds, but the discovery of VMware servers fails.
`Unable to obtain the list of trusted Certificate Authorities.`	This error might occur when VMware server credentials are added, updated, or validated. It occurs if the environment is configured to enabled communication between NetBackup (primary server, media server, or client) and vCenter, ESX, or any other VMware entity using authenticated certificates. Workaround: Ensure that certificates are installed and are valid.

Troubleshooting the browsing of VMware servers

The following table describes the problems that may occur when you click on a server under **VMware servers**.

Table 23-8 Errors browsing VMware servers

Issue	Explanation
No VMs or other objects were discovered for the VMware server.	■ If the server was added recently, the VM discovery process for that server may not have completed yet. Recommended action: Wait for the discovery process to finish. **Note:** The discovery of VMs and other objects in the vCenter, ESXi server, or VMware Cloud Director server begins: when server credentials are added or updated through the web UI or an API. However, the server's VMs and other objects might not appear in the UI immediately. They appear after the discovery process for the VMware server completes. Discovery also occurs at set intervals according to the `VMWARE_AUTODISCOVERY_INTERVAL` option. (The default interval is every 8 hours.) To perform autodiscovery of VMware server objects at a different frequency: See "Change the autodiscovery frequency of VMware assets" on page 65. ■ VMs or other objects of the VMware server may not be accessible for the added VMware server credentials. Recommended action: From the option menu on the right of the row, select **Edit**. Review the VMware server credentials and correct them as needed.

Troubleshooting the status for a newly discovered VM

The following table describes a problem that may occur when you review the status of a newly discovered VM under **Virtual machines**.

Table 23-9 Errors encountered when you review Status for a newly discovered VM

Error message or cause	Explanation and recommended action
The protection status of a VM indicates that it has not been backed up. However, a backup job that includes the VM has successfully completed.	In the NetBackup web UI, the protection status for a newly discovered VM does not indicate that it is backed up until the next backup of the VM has completed. In some circumstances, a new VM is backed up before the discovery of that VM has happened, as in the following scenario: ■ By default, autodiscovery occurs every 8 hours. ■ A new VM is added to the environment. ■ A backup job completes successfully before discovery completes. For example, a backup job that uses existing policies where the new VM is included as part of the backup selection criteria. ■ Later, discovery completes. However, in the NetBackup web UI, the protection status of the VM indicates that it has not been backed up. If you encounter a similar situation, you can still browse the recovery points and recover them. However, it is only after another backup of the VM successfully completes that the protection status indicates that the VM has been backed up. To review the protection status of a newly discovered VM in the NetBackup web UI, Cohesity recommends that you wait until the next successful backup has completed. Then, the protection status of the VM should correctly indicate its protection status.

Troubleshooting policy configuration

The following table describes the issues that may occur when you configure VMware policies.

Table 23-10 Errors with VMware backups

Issue	Explanation
When you select virtual machines on the policy **Clients** tab, NetBackup cannot obtain the host name, IP address, or DNS name of the virtual machine. Policy validation may also fail.	Make sure that the VMware Tools are installed and up to date on each virtual machine.

Table 23-10 Errors with VMware backups *(continued)*

Issue	Explanation
The Test query operation fails for a VMware policy.	There is no direct connectivity between the NetBackup primary server and the ESX server (for example, ESX server in NAT environment) where a virtual machine is to be backed up. Also, the **NetBackup host to perform automatic virtual machine selection** option is set to **Backup media server** (for example, NAT media server).

Troubleshooting the download of files from an instant access VM

The following table describes the problems that may occur when you download individual files from an instant access VM.

Table 23-11 Errors in downloading files

Issue	Explanation

Chrome: `This site can't be reached`
Firefox: `Server not found`
Edge: `Hmmm…can't reach this page`
	The web UI is unable to access the NetBackup media server with the name or IP address that the NetBackup primary server uses to connect to that media server.
	For example: If the primary server connects to the media server using `MSserver1.veritas.com`, the web UI must also be able to reach `MSserver1.veritas.com`. If the primary server uses a short name for the media server such as `MSserver1`, the web UI must be able to reach `https://MSserver1/...`
	Recommended action: Verify that the primary server and the web UI use the same name or IP address to access the media server (check the `hosts` file). For example: If the primary server uses the media server's short name, add the media server's short name and IP address to the `hosts` file of the PC or other host where the web UI is running.
	The hosts file location on Windows:
	`C:\Windows\System32\drivers\etc\hosts`
	The hosts file location on UNIX or Linux:
	`/etc/hosts`

	The web UI is unable to access the NetBackup media server because that server is behind a firewall. **Recommended action:** Contact the NetBackup security administrator.

Troubleshooting backups and restores of excluded virtual disks

Refer to the following table if you encounter restore issues for a backup that was configured to exclude virtual disks.

Table 23-12 Issues with excluding virtual disks

Issue	Explanation
The boot disk was backed up even though it was excluded from the backup.	The virtual machine only has a boot disk and no other disks.
	The boot disk is part of a managed volume (Windows LDM or Linux LVM). NetBackup can only exclude a boot disk if it is fully contained on a single disk.
	The virtual machine's boot disk is an independent disk and has no other disks.
	NetBackup was not able to identify the boot disk. The boot disk must include the boot partition and the system or the boot directory.
A restored boot disk has no data.	The boot disk is an independent disk. NetBackup cannot back up the data in this type of disk.
A restored virtual machine has a disk that contains missing or incomplete data.	The disk that has missing or incomplete data was excluded from the backup.
A data disk (or disks) was backed up even though it was excluded from the backup.	The virtual machine has only one disk (such as C:). In this case, the single drive is backed up and is not excluded.
A virtual machine is restored to an unexpected state.	You added a disk to the virtual machine and changed the settings that exclude disks. However, you did not create a backup of the entire virtual machine after you made the change.
Not all files can be restored individually.	If you remove disks from the custom attribute value between the differential backups, only those files that changed since the last backup can be restored individually. Alternatively, you can restore the entire virtual disk or the VM. After the next full backup, you can restore any of the files individually.
	If you remove controllers from **Specific disks to be excluded** between the differential backups, only those files that changed since the last backup are available for restore. All files are available for restore after the next full backup.
If you remove a disk from exclusion, the individual files that were last modified before the most recent backup cannot be restored.	To restore those files, restore the entire virtual disk or the virtual machine. After the next full backup, those files are available to restore individually.

How to determine the ESX network that NetBackup used for the backup or restore

If a virtual machine's disks are accessible to multiple ESX hosts, the disks can be accessed through any of the ESX hosts. The ESX host that is used for the access may or may not be the ESX host where the virtual machine is running or registered. All of the following must be accessible to each other and should have DNS configured:

- The vCenter server.
- All ESX hosts under the vCenter that have access to the virtual machine's vmdk files.
- The backup host.

If all hosts are not accessible to each other, the backup or restore may not succeed. In that case, you must determine which network NetBackup used for the backup or restore.

Note: For an NBD transport mode backup through vCenter, NetBackup uses the ESX network over which the ESX host was added or registered to the vCenter. For an NBD transport mode backup directly from the ESX host, NetBackup uses the ESX host's DNS/IP network.

The VxMS provider logs contain information on the network that NetBackup used.

See "Configuring VxMS logging" on page 319.

Check the VxMS provider logs for messages similar to those in this example:

```
10:49:21.0926 : g_vixInterfaceLogger:libvix.cpp:1811 <INFO> : Opening file
[MYDATASTORE] TestVM/TestVM-000001.vmdk (vpxa-nfc://[MYDATASTORE]
TestVM/TestVM-000001.vmdk@MyESX.xxx.xxx.com:902)

   10:49:22.0301 : g_vixInterfaceLogger:libvix.cpp:1811 <INFO> : DISKLIB-LINK  :
Opened 'vpxa-nfc://[MYDATASTORE]
TestVM/TestVM-000001.vmdk@MyESX.xxx.xxx.com:902' (0x1e): custom, 41943040
sectors / 20 GB.

   10:49:22.0301 : g_vixInterfaceLogger:libvix.cpp:1811 <INFO> : DISKLIB-LIB   :
Opened "vpxa-nfc://[MYDATASTORE]
TestVM/TestVM-000001.vmdk@MyESX.xxx.xxx.com:902" (flags 0x1e, type custom).

   10:49:22.0301 : vdOpen:VixInterface.cpp:480 <DEBUG> : Done with
VixDiskLib_Open(): 200346144
   10:49:22.0301 : openLeafSnapshotDisks:VixGuest.cpp:475 <DEBUG> : vdOpen()
```

```
succeess
   10:49:22.0301 : openLeafSnapshotDisks:VixGuest.cpp:476 <INFO> : Transport
mode in effect = nbd
```

VMware logs the messages starting with `g_vixInterfaceLogger`. Such messages in the example indicate that TestVM-000001.vmdk is opened over the ESX host network `MyESX.xxx.xxx.com`.

The following article contains related information:

- *Best practices when using advanced transport for backup and restore* http://kb.vmware.com/kb/1035096

Preventing browsing delays caused by DNS problems

NetBackup may be unable to identify virtual machines when you use the **Browse for Virtual Machines** dialog. Virtual machine host names may not be properly configured in your Domain Name Server system (DNS), or the DNS system may be slow. A timeout message may appear, or messages similar to the following may appear in the NetBackup detailed status log:

```
17:25:37.319 [12452.10360] get_vSphere_VMs: Processing vm 002-wcms
17:25:37.319 [12452.10360] get_vSphere_VMs:    uuid
421a8b46-063d-f4bd-e674-9ad3707ee036
17:25:37.319 [12452.10360] get_vSphere_VMs:    vmxdir [san-05] 002-wcms/
17:25:37.319 [12452.10360] get_vSphere_VMs:    datastore san-05
17:25:37.319 [12452.10360] get_vSphere_VMs:    IpAddress 172.15.6.133
17:25:41.866 [12452.10360] get_vSphere_VMs: retry_gethostbyaddr for
172.15.6.133 failed with The requested name is valid, but no data of
the requested type was found.
```

Note: NetBackup may be unable to determine the host names of the virtual machines from their IP addresses (reverse lookup may fail).

To prevent browsing delays caused by DNS problems (Windows)

1. On the Windows desktop of the backup host, click **Start > Run** and enter `regedit`.
2. To be on the safe side, make a backup of the current registry (**File > Export**).
3. Go to **HKEY_LOCAL_MACHINE > SOFTWARE > Veritas > NetBackup > CurrentVersion > Config** and create a key that is called `BACKUP`.

4 Create a new DWORD under `BACKUP`, called `disableIPResolution`.

This registry key causes NetBackup to use the virtual machine's IP address as the virtual machine's host name.

5 Use the NetBackup **Browse for Virtual Machines** screen to rediscover the virtual machines. The host names should now be the IP addresses.

See "Browse for VMware virtual machines" on page 107.

To prevent browsing delays caused by DNS problems (Linux)

1 On the Linux backup host, create (or open) the following file:

```
/usr/openv/netbackup/virtualization.conf
```

2 Add the following to the file:

```
[BACKUP]
"disableIPResolution"=dword:00000000
```

This entry causes NetBackup to use the virtual machine's IP address as the virtual machine's host name.

Note: If the file already contains a `[BACKUP]` line, do not add another `[BACKUP]` line. Any other lines that already exist under `[BACKUP]` should remain as they are.

3 Use the NetBackup **Browse for Virtual Machines** screen to rediscover the virtual machines. The host names should now be the IP addresses.

See "Browse for VMware virtual machines" on page 107.

The following applies if: the **Primary Identifier** in VMware Intelligent Policies is selected as **VM host name** and **Reverse name lookup** is enabled in the configuration setting.

In a large VMware environment, reverse name lookups can be very slow depending on the number of virtual machines being discovered. You can change the `VNET_OPTIONS` option to determine how many items NetBackup can cache. This value is in the `bp.conf` file on UNIX and Linux, and the registry on Windows.

The third value 200 is the default number of entries to be cached. Each entry takes about 1 kilobyte in memory. Available memory needs to be taken into account when you change this value. The maximum number of allowed entries is 100000.

```
VNET_OPTIONS = 120 3600 200 40 3 1 30 10 1793 32 0 0
```

Use the `nbgetconfig` command to view the configuration settings. Use `nbsetconfig` to change the settings.

Changing the browsing timeout for virtual machine discovery

You can increase the browsing timeout value to improvement performance when you browse for virtual machines.

To change the browsing timeout value

1. On the host that runs the NetBackup Administration Console, open the following file:

   `/usr/openv/java/nbj.conf`

2. Change the value of the `NBJAVA_CORBA_DEFAULT_TIMEOUT` parameter.

   By default, this parameter is set to 60 seconds:

   `NBJAVA_CORBA_DEFAULT_TIMEOUT=60`

   Increase the value to a higher number.

   For more information on the `nbj.conf` file, see the NetBackup Administrator's Guide, Volume I.

See "Browse for VMware virtual machines" on page 107.

Changing timeout and logging values for vSphere

Table 23-13 lists the vSphere keys and their default values for various timeouts. These values can be changed on the backup host (see the procedure in this topic).

Table 23-13 DWORD keys and defaults for vSphere timeouts

DWORD key name	Default value (in seconds)
`jobtimeout`	900
`poweroptimeout`	900
`snapshottimeout`	900
`registertimeout`	180
`browsetimeout`	180
`connecttimeout`	300

The key and default for the vSphere API logging level are the following.

Table 23-14 DWORD key and default for vSphere API log level

DWORD key name	Default value
`vmcloglevel`	0 (no logging)

Changes to the vSphere API logging level affect the following logs on the backup host:

- For backups (snapshot creation): `bpfis` log
- For restores: `bpVMutil` log
- For virtual machine discovery: `ncfnbcs` log (originator ID 366)

The logs are in the following location on the backup host:

Windows: *install_path*\NetBackup\logs\

Linux: /usr/openv/netbackup/logs

To change vSphere timeouts and logging values on Windows

1. On the Windows desktop of the backup host, click **Start > Run** and enter `regedit`.
2. To be on the safe side, make a backup of the current registry (**File > Export**).
3. Go to **HKEY_LOCAL_MACHINE > SOFTWARE > Veritas > NetBackup > CurrentVersion > CONFIG** and create a key that is called `BACKUP`.
4. To change a timeout value, create a new DWORD under `BACKUP`, using the appropriate registry name (such as `jobtimeout` or `poweroptimeout`).

   Enter a value for the timeout.
5. To change the level of vSphere API logging, create a new DWORD called `vmcloglevel` and enter the new logging value.

   The allowed values are 0 through 6, where 0 is no logging and 6 is the highest log level.

To change vSphere timeouts and logging values on Linux

1. On the Linux backup host, create (or open) the following file:

   `/usr/openv/netbackup/virtualization.conf`

2. To change a timeout value, enter a new `dword` line under `[BACKUP]`, using the appropriate name (such as `jobtimeout` or `poweroptimeout`). Include a value for the timeout.

   For example:

   ```
   [BACKUP]
   "jobtimeout"=dword:60
   ```

   This example sets the job timeout to 60 seconds.

   Note: If the file already contains a `[BACKUP]` line, do not add another `[BACKUP]` line. Any other lines that already exist under `[BACKUP]` should remain as they are.

3. To change the level of vSphere API logging, enter a `dword` line for `vmcloglevel` with a logging value, under `[BACKUP]`. For example:

   ```
   "vmcloglevel"=dword:6
   ```

   The allowed values are 0 through 6, where 0 is no logging and 6 is the highest log level.

Credentials for VMware server are not valid

A number of issues can prevent NetBackup from gaining access to the ESX server or vCenter server. When you add credentials for a VMware server, NetBackup validates the credentials.

Problems can result for a variety of reasons, including the following:

- Problems with the host name or host IP address.
  - An incorrect virtual machine server name. Make sure that the server name is entered correctly.
  - You used a short host name for the server for its NetBackup credentials, and need to replace that name with the fully qualified host name.
  - Two hosts currently resolve to the same IP address, and one of them must be renamed and assigned a new IP address.

  For these situations, delete and re-add the VMware server.
  See "Remove VMware servers" on page 72.

See “Add VMware servers” on page 66.

- An invalid username or password. Make sure that a valid username and password were entered correctly.
- An incorrect port number. Make sure that the port number is correct in the VMware server credential settings. If the VMware server uses the default port, no port specification is required.
- You do not have enough privileges to perform backups or restores. (Note however that lack of sufficient privileges may not cause the credential validation to fail.)

  For the minimum permissions needed to back up and restore with vStorage, see the following tech note:

  https://www.veritas.com/docs/100001960

Snapshot error encountered (status code 156)

The following table describes the VMware issues that relate to NetBackup status code 156.

Table 23-15 Possible causes of status code 156

Causes of status code 156	Description and recommended action
NetBackup cannot obtain the volume ID of a drive	NetBackup may not be able to obtain the volume ID of a drive. In that case, none of the virtual machine drives are backed up. The backup fails with NetBackup status code 156. The drive may be down.
A backup of the virtual machine is already active	You cannot run more than one backup per virtual machine at a time. If you start a second backup of the virtual machine while the first backup is active, the second job fails with a status 156. Recommended action: Wait until the first job completes, then run the second one.

Table 23-15 Possible causes of status code 156 *(continued)*

Causes of status code 156	Description and recommended action
Cannot find virtual machine name	NetBackup cannot find the host name or VM display name of a virtual machine that is listed in the backup policy. The detailed status log may include the following error message: `Critical bpbrm (pid=<pid number>) from client <client name>: FTL - snapshot creation failed, status 156.)` If the virtual machines do not have static IP addresses, you can configure NetBackup to identify virtual machines by their VM display names or UUIDs. Examples of the environments that do not use static IP addresses are clusters, and the networks that assign IP addresses dynamically. Note that NetBackup may have been configured to identify virtual machines by their VM display names. In that case, make sure that the display names are unique and that they do not contain special characters. See "Primary VM identifier options (VMware)" on page 93.
The virtual machine is powered off	Through a vCenter server, NetBackup can back up the virtual machines that are turned off. You must provide credentials for NetBackup to access the vCenter server. See "Add VMware servers" on page 66. If NetBackup uses credentials for an ESX server instead of vCenter, it may not be able to identify a turned off virtual machine. Note the following: ■ If the policy uses VM host name or VM DNS name as the Primary VM identifier, NetBackup may not find the virtual machine. The backup fails. ■ If the policy uses VM display name or VM UUID as the Primary VM identifier, NetBackup can identify the virtual machine. The backup succeeds.
The virtual machine has one or more independent disks and is in a suspended state	If a virtual machine with independent disks is in a suspended state, snapshot jobs fail. Messages similar to the following appear in the job details log: `01/12/2015 17:11:37 - Critical bpbrm (pid=10144) from client <client name>: FTL - VMware error received: Cannot take a memory snapshot, since the virtual machine is configured with independent disks.` This issue results from a VMware limitation (SR#15583458301). More information is available in the following VMware article: http://kb.vmware.com/kb/1007532 As a workaround, change the state of the virtual machine to powered on or powered off, and rerun the backup. **Note:** Data on independent disks cannot be captured with a snapshot. The rest of the virtual machine data is backed up.

Table 23-15 Possible causes of status code 156 *(continued)*

Causes of status code 156	Description and recommended action
The virtual machine's disk is in raw mode (RDM)	The RDM is ignored (not backed up) and any independent disk is recreated but empty. See "Configurations for backing up RDMs" on page 375.
The attempt to create a snapshot exceeded the VMware timeout	If the attempt to create a snapshot of the virtual machine exceeds the VMware timeout of 10 seconds, the snapshot fails with NetBackup status 156. This timeout may occur if the virtual machine is configured with a large number of volumes. Note that the timeout may be encountered even if the **Virtual machine quiesce** option was disabled. Do one of the following: ■ Reduce the number of volumes within the virtual machine. ■ Install a NetBackup client on the virtual machine and select another backup method for the policy (not the VMware snapshot method).
The virtual machine has no vmdk file assigned	The snapshot fails if the virtual machine has no vmdk file. Virtual machines without vmdk files can occur in a vCenter Site Recovery Manager (SRM) environment. If a replicated virtual machine has never been active, it is in passive mode and may have no vmdk files. You can enable the **Ignore diskless VMs** option on the **VMware Advanced Attributes** tab of the policy. If this option is enabled: NetBackup does not back up a replicated (passive) virtual machine in an SRM environment if that virtual machine has no vmdk files. More information is available on the **Ignore diskless VMs** option. See "VMware - Advanced attributes" on page 96.

Table 23-15 Possible causes of status code 156 *(continued)*

Causes of status code 156	Description and recommended action

The vmdk file has too many delta files	Whenever a VMware snapshot occurs, a delta.vmdk file is created for each vmdk. If 32 or more such delta files exist for a single vmdk file, a NetBackup backup of that VM may fail (status 156). The NetBackup Activity Monitor job details contain messages similar to the following:
	`02/06/2015 10:33:17 - Critical bpbrm (pid=15799) from client fl5vm1_2012: FTL - vSphere_freeze: Unable to proceed with snapshot creation, too many existing delta files(44).`
	`02/06/2015 10:33:17 - Critical bpbrm (pid=15799) from client fl5vm1_2012: FTL - VMware_freeze: VIXAPI freeze (VMware snapshot) failed with 25: SYM_VMC_FAILED_TO_CREATE_SNAPSHOT`
	`02/06/2015 10:33:17 - Critical bpbrm (pid=15799) from client fl5vm1_2012: FTL - vfm_freeze: method: VMware_v2, type: FIM, function: VMware_v2_freeze`
	To back up the VM, do the following:
	1 Consolidate the VM's snapshots.
	In the VMware interface, right-click on the VM and select **Snapshot > Consolidate**. For more information, see your VMware documentation.
	2 Verify that each of the VM's vmdk files now has fewer than 32 delta files.
	If the snapshot consolidation was not successful, see the following VMware article for further assistance:
	Committing snapshosts in vSphere
	3 Rerun the NetBackup backup.

VMware snapshot quiesce operation failed	If the NetBackup policy is enabled for virtual machine quiesce (the default), the VMware snapshot operation in vSphere initiates a quiesce of the virtual machine. If the snapshot quiesce fails, the NetBackup job fails with status 156.

The origin of the snapshot failure: NetBackup or VMware?

When a NetBackup snapshot job fails with status 156, the problem may originate in your VMware environment rather than in NetBackup. You can begin to isolate the problem to one environment or the other by using vSphere Client to take a snapshot of the VM. NetBackup support often uses this approach to investigate a snapshot issue.

To identify the environment in which the snapshot error occurred

1. In the vSphere interface, right-click on the VM and click **Snapshots > Take Snapshot**.
2. In the **Take VM Snapshot for** dialog, click **Quiesce guest file system** if the NetBackup policy was enabled for virtual machine quiesce (the default).

   In the NetBackup policy, the **Virtual machine quiesce** option is in the **VMware - advanced attributes** on the **VMware** tab.

   Note: In the **Take VM Snapshot for** dialog, make sure the **Snapshot the virtual machine's memory** option is not selected. NetBackup does not use that option.

3. Start the snapshot and check the **Recent Tasks** pane for snapshot status.
   - If the snapshot does not complete, the problem with the NetBackup snapshot may be in the VMware environment. Consult your VMware documentation.
   - If the VMware snapshot is successful, the issue may be with NetBackup. For relevant error messages, consult the NetBackup bpfis logs.
     See "NetBackup logging for VMware" on page 316.
     The following topic summarizes some common causes of 156 errors:
     See " Snapshot error encountered (status code 156)" on page 340.

Conflict between NetBackup and VMware Storage vMotion with vSphere 5.0 or later

To avoid conflicts with Storage vMotion in vSphere 5.0 or later, NetBackup should conduct backups through the vCenter server, not through the ESX host. A backup directly through the ESX server may fail if Storage vMotion simultaneously migrates the virtual machine's files. In addition, the virtual machine's snapshot files may be stranded or other problems with the virtual machine may result. VMware has acknowledged this issue.

If the backup fails, the NetBackup job details contain a message similar to the following:

```
Error opening the snapshot disks using given transport mode: Status 23.
```

To back up a virtual machine while its files are in the process of migration, NetBackup must conduct the backup through the vCenter server.

To back up a virtual machine while its files are in the process of migration

1. Open the NetBackup web UI.
2. On the left, select **Workloads > VMware**. Select the **VMware servers** tab and select **Add**.
3. Select **vCenter**.
4. Provide the other details and the credentials for the server.
5. Select **Save**.
6. Locate and select the **ESXi server**. Then select **Delete**.
7. On the **VMware servers** tab, select the **Add** button.
8. Select **Restore ESXi**.
9. Provide the other details and the credentials for the server.
10. Rerun the backup.

Backup or restore job hangs

NetBackup may have exceeded the maximum number of allowed VMware NFC connections to the ESX server when it used the transport modes `nbd` or `nbdssl`. Note that NetBackup uses one NFC connection per virtual disk on the ESX or ESXi server.

If NetBackup is configured to access the ESX server directly (not through a vCenter or VirtualCenter server), fewer connections are allowed. The following are the maximum connections as set by VMware:

Table 23-16 VMware NFC connection limits for `nbd` or `nbdssl` transfers

ESX version	Type of access to the ESX server	Maximum NFC connections allowed
ESX 4	Directly to ESX server	9
ESX 4	Through vCenter	27
ESXi 4	Directly to ESX server	11
ESXi 4	Through vCenter	23
ESXi 5	Directly to ESX server	The maximum total for all NFC connection buffers to an ESXi host is 32 MB
ESXi 5	Through vCenter	52

Try a different transport type (such as SAN or hotadd). If a different transport type is not available and NetBackup accesses the ESX servers directly, set up access through a vCenter (or VirtualCenter) server. Use of a server increases the maximum number of allowed connections. For example: With 27 connections, NetBackup can access a virtual machine that has up to 27 disks, if the ESX 4 server is behind a vCenter server.

Note that the connection limits are per-host (that is, per vCenter or ESX server).

For example, assume the following environment:

- An ESX 4.0 server with three virtual machines.
- Each virtual machine has ten virtual disks.
- The virtual machines are behind a vCenter 4.0 server.

For a simultaneous backup of the three virtual machines, NetBackup requires 30 NFC connections. With a limit of 27 NFC connections per vCenter server, any of the three backup jobs may hang.

These limits are described in the *VMware Virtual Disk API Programming Guide*:

VMware Virtual Disk API Programming Guide

See also the following section of the *VMware vSphere 5 Documentation Center*:

Virtual Disk Transport Methods

VMware SCSI requirements for application quiesce on Windows

For a snapshot that quiesces an application on a Windows VM, VMware imposes the following disk requirements:

- The VM's disks must be SCSI, not IDE.
- The SCSI disks on the VM's SCSI controller must not occupy more than half of the total number of slots in the controller. Since the controller has a total of 15 slots, the number of disks in that controller must not exceed 7.
  If these conditions are not met and the NetBackup VMware policy enables the **Virtual machine quiesce** option, the backup may fail with status code 156.
  For more information on this VMware requirement, refer to the information on Windows Backup Implementations on the VMware docs web site.
  https://docs.vmware.com/

VMware virtual machine does not restart after restore

The virtual machine may have been configured as follows:

- At the time of the backup, the virtual machine had a combination of SATA and SCSI disks, or of SATA, SCSI, and IDE disks.
- The guest OS resided on one of the SCSI disks.

The virtual machine when restored may attempt to boot from the SATA or the IDE disk. The boot attempt fails with the message "Operating system not found."

VMware has identified this problem and will address it in a future release.

As a workaround, reconfigure the BIOS on the virtual machine to boot from the correct SCSI disk.

A restored VM may not start or its file system(s) may not be accessible

A restored VM may not boot up or its file system(s) may not be accessible in the following case:

- The VM's guest operating system is Windows 8,
- The VM is restored from a block-level incremental backup image,
- And the restore uses the hotadd transport mode.

As a result of a VMware issue in VDDK 5.5.x, the Windows NTFS Primary File Table on the restored VM may be corrupted. As a workaround, use a different transport mode to restore the VM (not hotadd).

NetBackup job fails due to update tasks on the VMware server

Certain virtual machine update tasks on the VMware server may cause a NetBackup job to fail. For example, a restore may fail when you use the nbrestorevm command to restore multiple VMs at the same time. If the failed restore job creates the VM, the VM is deleted.

On the backup or restore host, the `bpVMutil` log may contain a message similar to the following:

```
Detail: <ManagedObjectNotFoundFault xmlns="urn:vim25"
xsi:type="ManagedObjectNotFound"><obj type="VirtualMachine">
vm-14355</obj>
        </ManagedObjectNotFoundFault>
```

As a workaround, rerun the job for the operation that failed.

Note: If possible, avoid VMware maintenance activities during backup or restore operations.

The vSphere interface reports that virtual machine consolidation is needed

When NetBackup begins a virtual machine backup, it requests a VMware snapshot of the virtual machine in vSphere. If the NetBackup policy is enabled for virtual machine quiesce (the default), the VMware snapshot operation initiates a quiesce of the virtual machine. If snapshot quiesce fails, the NetBackup job fails with status 156 and VMware snapshot delta files may be left behind in vSphere. Note: As a result of leftover snapshot delta files, the vSphere status for the virtual machine may warn that virtual machine consolidation is needed.

VMware has acknowledged the problem of leftover delta files after a snapshot quiesce failure. VMware has fixed this issue in certain ESXi versions. See the following VMware article for more information on this issue:

Delta disk files (REDO logs) are left uncommitted after a failed quiesced snapshot operation (2045116)

For a workaround for leftover delta files, see the following VMware article:

Committing snapshots when there are no snapshot entries in the Snapshot Manager (1002310)

Linux VMs and persistent device naming

For Linux VMs without persistent device naming, multiple disk controllers (such as IDE, SCSI, and SATA) may complicate the recovery of individual files. This issue occurs because non-persistent device naming, such as `/dev/sda` and `/dev/sdb`, may cause unexpected mount point changes after a restart. If the VM has a SCSI disk and SATA disk, the Backup, Archive, and Restore interface may show incorrect mount points for the VM's files. For example, the files originally under `/vol_a` might appear under `/vol_b` when you browse to restore them. The restore is successful, but the restored files may not be in their original directories.

As a workaround, search for the files on the restored VM and move them to the proper locations.

To prevent this issue on Linux VMs with multiple disk controllers, it is recommended a persistent device-naming method for mounting the file systems. When persistent naming is in place, device mounting is consistent and this issue does not occur when you restore files from future backups.

For persistent device naming, you can mount devices by UUIDs. The following is an example of the `/etc/fstab` file that contains the devices that are mounted by UUIDs:

```
UUID=93a21fe4-4c55-4e5a-8124-1e2e1460fece /boot  ext4  defaults  1 2
UUID=55a24fe3-4c55-4e6a-8124-1e2e1460fadf /vola  ext3  defaults  0 0
```

Note: Limit the number of characters for each `fstab` entry to 90 on a VMware VM.

To find the device UUIDs, you can use either of the following commands:

```
blkid
```

```
ls -l /dev/disk/by-uuid/
```

Note: NetBackup also supports the by-LABEL method for persistent device naming.

For a VMware virtual machine with Windows dynamic disks, a restore from incremental backup fails with a Windows restore host and the hotadd transport mode

A restore of a Windows virtual machine by means of the hotadd transfer mode may fail in the following case:

- A backup is taken of a Windows virtual machine that has a dynamic disk group.
- After the backup, another dynamic disk is added to the virtual machine's disk group.
- After the dynamic disk is added, an incremental backup is taken of the virtual machine.
- A Windows restore host is used with the hotadd transport mode to restore the virtual machine from the incremental backup.

The restore fails when the Windows restore host tries to mount the dynamic disk that was added after the first backup. Depending on the data that has already been restored, Windows may detect the dynamic disk as `Invalid` or `Foreign`. Further writes to an `Invalid` or `Foreign` disk are unsuccessful and the restore fails.

The restore fails with status 1, "the requested operation was partially successful." Messages similar to the following may appear in the VxMS provider logs:

```
14:10:18.0854 : vdWrite:../VixInterface.cpp:760 <ERROR> : Error
24488361628532739 in write with vdhandle 48870608 startsector
128 numsectors 1 14:10:18.0854 : write:VixFile.h:333 <ERROR>
: Returned error 3, offset 0x0000000000010000, length
0x0000000000000200   14:10:18.0854 : write:VixFile.h:334
<ERROR> : Returning: 11
 14:10:18.0854 : vixMapWrite:../VixCoordinator.cpp:1259 <ERROR>
: Returning: 11
14:10:18.0854 : vix_map_write:../libvix.cpp:1826 <ERROR>
: Returning: 11
```

When they are enabled, VxMS logs are written in the following directory:

Windows:

install_path\NetBackup\logs\vxms

Linux:

/usr/openv/netbackup/logs/vxms

Note: For successful restores from future incremental backups, run backups with the **Use Accelerator** option in the policy.

Try any of the following workarounds to restore from the current incremental backup:

- Use a Linux restore host (not Windows).
- Use a different transport mode, such as NBD, NBDSSL, or SAN (not hotadd).
- When the dynamic disk (the one that was added after the first backup) is mounted for restore, manually set the disk to offline. When the disk is offline, NetBackup can write data to it and successfully complete the restore.
  See the remainder of this tech note for assistance with this workaround.

To determine when the dynamic disk is mounted for restore

- Use the Windows Disk Management utility (**Control Panel > Administrative Tools > Computer Management > Disk Management**), or run `diskpart` in administrator mode and enter the `list disk` option.

  When Windows attempts to mount the disk, it labels the disk as `Invalid` or `Foreign`.

To use diskpart to take the dynamic disk offline

1. On the Windows restore host, run `diskpart` in administrator mode.
2. Enter `list disk` to list all disks and find the `Invalid` or `Foreign` disk.
3. Enter `select disk` *`disk ###`* to select the disk that is `Invalid` or `Foreign`.
4. Enter `offline disk` to offline the disk.

   Example session:

```
DISKPART> list disk

  Disk ###  Status         Size     Free     Dyn  Gpt
  --------  -------------  -------  -------  ---  ---
  Disk 0    Online          100 GB  1024 KB
  Disk 1    Online          256 GB    56 GB        *
  Disk 2    Invalid          40 MB    40 MB   *    *
  Disk 3    Offline          40 MB    40 MB

DISKPART> select disk 2

Disk 2 is now the selected disk.

DISKPART> offline disk

DiskPart successfully taken offline the selected disk.
```

Simultaneous hotadd backups (from the same VMware backup host) fail with status 13

During simultaneous backups from the same VMware backup host, some of the backups may fail with status 13, "file read failed." A hotadd backup of multiple disks may take more time than the client-read timeout allows (the default is 300 seconds). The delay may be caused by locking timeouts in the VMware VDDK.

In the NetBackup Activity monitor, the detailed status log may include messages similar to the following:

```
12/05/2014 06:43:53 - begin writing
12/05/2014 06:48:53 - Error bpbrm (pid=2605) socket read failed:
errno = 62 - Timer expired
12/05/2014 06:48:55 - Error bptm (pid=2654) media manager terminated
by parent process
```

The `/NetBackup/logs/vxms` log may include repeated instances of a VDDK message similar to the following:

```
12/08/2014 05:11:35 : g_vixInterfaceLogger:libvix.cpp:1844 <DEBUG> :
[VFM_ESINFO] 2014-12-08T05:11:35.146-06:00 [7F1B1163F700 info Libs']
FILE: FileLockWaitForPossession timeout on '/var/log/vmware/hotAddLock.
dat.lck/M34709.lck' due to a local process '15882-26732358(bpbkarv)'
```

To prevent this issue, do one of the following:

- Reduce the number of hotadd backups that run simultaneously.
- Increase the client-read timeout on the media server as appropriate (15 minutes or more):
  In the NetBackup web UI, click **Hosts > Host properties**. Select the media server. Click **Connect**. Click **Timeouts > Client read timeout**.

Troubleshooting VMware tag usage

Tag associations are backed up and restored as part of the VMware backup process. Backup and restore of tag associations is a best effort. Any tag collection errors are shown in the Activity Monitor for the virtual machine snapshot job.

Because of the best effort on backup and restore of tag associations, you may receive unexpected behavior. Please be aware of the following:

- You can receive a NetBackup Status Code 0 even if the tag associations are not captured during the backup. Because tag backup is best effort, this error is not considered a failure that halts a backup. Any tag collection errors that occur in the backup are shown in the Activity Monitor for the virtual machine snapshot job.
- You can receive a NetBackup Status Code 0 on a restore even if tag associations are not restored. Any tag collection errors that occur in the backup are shown in the Activity Monitor for the virtual machine snapshot job.
- Restores receive a NetBackup Status Code 1 when:

- The tag doesn't exist on the target vCenter Server. Be aware that NetBackup restores tag associations to virtual machines by the tag identifier not tag name.
- The virtual machine was restored to a pre-6.0 vCenter Server.
- Other VMware failures.

Query behavior with unsupported versions of VMware

Table 23-17 NetBackup query behavior for tag field keyword with unsupported versions of VMware

Version of VMware	Behavior	Additional information
vCenter earlier than 5.1	Virtual machines are evaluated as if tags are not set.	N/A
vCenter 5.1 and 5.5	Virtual machine selection is marked as **Failed**.	If tag information is required to determine virtual machine selection, the virtual machine is marked as **Failed** with NetBackup Status Code 4266.
ESX server	Virtual machine selection is marked as **Failed**.	If tag information is required to determine virtual machine selection, the virtual machine is marked as **Failed** with NetBackup Status Code 4265.

List of possible Tag values not displayed in Query Builder

If the list of possible values does not return any results when the **Tag Field** is selected in the **Policy Query Builder**, there are several causes of this problem.

- Confirm that your version of vCenter Server is 6.0 or later.
  NetBackup support for tags starts with VMware vCenter Server 6.0.
  See "Notes and limitations for tag usage in VMware Intelligent Policy queries" on page 46.
- Confirm there are tags defined using the vSphere Web Client.
- Confirm that all tags and categories have descriptions.
  If one of the tags does not have a description, when you attempt to browse for tags in the VMware Intelligent Policy, no tags are displayed.

Cohesity has confirmed that this issue is resolved in vCenter Server 6.0 Update 1.

VMware Knowledge Base article: http://kb.vmware.com/kb/2124204

- Confirm the system times of the discovery host and the vCenter Server are synchronized.
  VMware Knowledge Base article: http://kb.vmware.com/kb/2125193

Timeout issues in discovery job

In large VMware environments, you may experience timeout issues during the discovery job or the test query. Review the NetBackup Activity monitor for this message:

```
09/21/2015 10:23:05 - Error nbpem (pid=13064) VMware vCloud Suite
SDK API Failed, msg = [This method requires authentication.], display
name = [display_name], server = [server_name]
```

This message indicates that the discovery job exceeded the bearer timeout. You need to increase this timeout for the job to complete successfully.

To adjust the timeout on the vCenter

1. Open the VMware vSphere Web Client.
2. Select **Administration > Single Sign-On > Configuration > Policies > Token Policy**.
3. Increase **Maximum bearer token lifetime** from the default 300 seconds. Because each environment is unique, Cohesity does not have any recommendations on this value. Increase the value until the problem is eliminated.

Ensuring that guest customizations can be restored in vCloud Director

NetBackup can back up VMware vCloud Director environments and restore virtual machines into vCloud Director.

See "About NetBackup for vCloud Director" on page 254.

To ensure that any VM guest customizations are restored into vCloud Director, you must set a NetBackup parameter, as follows:

- On Windows, you must set a registry value.
  See "To ensure that guest customizations can be restored in vCloud Director on Windows" on page 355.
- On UNIX and Linux, you must edit a NetBackup configuration file.

See "To ensure that guest customizations can be restored in vCloud Director on Linux" on page 355.

The parameter value specifies a wait period in seconds so that the guest customizations can be restored successfully. (The VMware API requires that the VMware Tools are installed and running, but the state of the VMware Tools cannot be identified after the restore. Therefore, we wait the specified amount of time so that the VMware Tools are running in the initial restore environment.)

To ensure that guest customizations can be restored in vCloud Director on Windows

1. On the Windows desktop of the backup host, click **Start > Run** and enter `regedit`.
2. To be on the safe side, make a backup of the current registry (**File > Export**).
3. Go to **HKEY_LOCAL_MACHINE > SOFTWARE > Veritas > NetBackup > CurrentVersion > CONFIG** and create a key that is called `BACKUP`.
4. Create a new DWORD under `BACKUP` and name it `powerCycleInterval`.

   Enter a decimal value of `60` for the timeout.
5. Close the Windows Registry Editor.

To ensure that guest customizations can be restored in vCloud Director on Linux

1. On the Linux backup host, create (or open) the following file:

   `/usr/openv/netbackup/virtualization.conf`
2. Create a line in the file named `[BACKUP]` and then on a separate line create a `powerCycleInterval dword` parameter with a value of `60`, as follows:

   For example:

   ```
   [BACKUP]
   "powerCycleInterval"=dword:60
   ```

   Note: If the file already contains a `[BACKUP]` line, do not add another `[BACKUP]` line. Any other lines that already exist under `[BACKUP]` should remain as they are.
3. Save the file and then close it from the text editor.

Troubleshooting vmdk restore to existing VM

For virtual disk restores and in-place restores to an existing VM, be aware of the following:

- If the VMDKs cannot be attached to the target VM, the restored VMDKs are retained on the temporary VM. The name of the temporary VM is available in the job details in NetBackup. In the following job details example, the temporary VM name is `vCenter60vm1_rhel6.4_1465584674`:

  ```
  06/10/2016 13:51:17 - Info bpVMutil (pid=3400) Restoring [datastore1]
   vCenter60vm1_rhel6.4/vCenter60vm1_rhel6.4_4.vmdk to [datastore1]
   vCenter60vm1_rhel6.4/vCenter60vm1_rhel6.4_4-1465584677.vmdk
  06/10/2016 13:51:38 - Info bpVMutil (pid=3400) Successfully created
   virtual machine vCenter60vm1_rhel6.4_1465584674 with specified disks.
  06/10/2016 13:51:41 - requesting resource @aaaab
  ```

  You can access the data on the temporary VM.
  If the temporary VM is retained after a restore failure, the restore job contains a message similar to the following:

  ```
  06/14/2016 15:29:06 - Info bpVMutil (pid=5225) attachDisksToExistingVM:
   Unable to attach restored disks to target VM vCenter60vm2_rhel6.4
  06/14/2016 15:29:06 - Info bpVMutil (pid=5225) attachDisksToExistingVM:
   Temporary VM  with restored virtual disks was left in place, it can be
   used to access restored data
  ```

  By default, NetBackup retains the temporary VM if the disks are not attached. To change that behavior, set the `DeleteRestoredVMOnError` field to `Yes` in the restore parameters file.

- After a restore to an existing VM, the next backup of the VM backs up the restored virtual disks. This backup may show a warning during collection of the Changed Block Tracking (CBT) information.

- For an In-place disk restore, raw devices (RDMs) and independent disks are not deleted or replaced during restore. If the controller for these disks conflicts with the disks being restored, the restore fails. The following example messages are job details from a failed in-place restore:

  ```
  May 07, 2020 10:26:21 AM - Warning bprd.sfr (pid=2425) Unable to
   attach the restored disks to requested VM
  May 07, 2020 10:26:21 AM - restored from image
   InPlaceDiskRestoreDemo_1588837243; restore time: 0:00:50
  May 07, 2020 10:26:21 AM - end Restore; elapsed time 0:00:50
  May 07, 2020 10:26:21 AM - Info bpVMutil (pid=2673)
   attachDisksToExistingVM: Controller scsi0-1 not available
   to perform in-place disk restore. Aborting restore.
   The requested operation was partially successful(1)
  ```

Troubleshooting backups of virtual machines on Virtual Volumes (VVols)

When troubleshooting backups of virtual machines on VVols, note the following:

- Each NetBackup snapshot job creates a vSphere snapshot of the virtual machine.
- To investigate snapshot failures, check the storage array's VASA provider logs as well as the vSphere error messages. (VASA is vSphere API for Storage Awareness.)

Table 23-18 Errors for backups of virtual machines on VVols

Error	Explanation
vSphere snapshot creation fails for a backup of a VM on VVol.	Ensure that you have the required snapshot license from the array vendor. Licensing requirements for vSphere snapshots vary from one type of VVol storage to another, depending on the array vendor. Each NetBackup snapshot job creates a vSphere snapshot of the virtual machine.
	There is insufficient space in VVol storage. Space requirements vary from one array vendor to another. Consult the storage array documentation.

Issues with the CA certificate during installation of the NetBackup client on VMware Cloud (VMC)

Failed to display CA certificate details

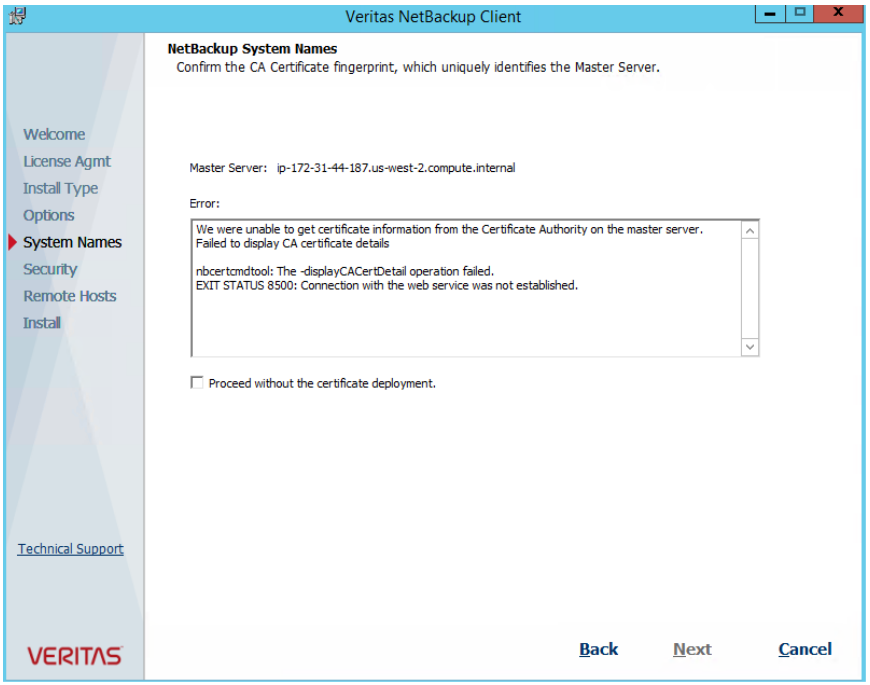

To fix this problem, do the following:

1. Click **Cancel** to cancel the installation.
2. Make sure the private DNS name was used to install the primary server.
3. In the hosts file of the primary server installed in AWS and of the backup host installed in VMC, add the following:
   - The private IP and private DNS name of the NetBackup primary server.
   - The IP and the DNS name of the backup host.
     The hosts file location on Windows:
     `C:\Windows\System32\drivers\etc\hosts`
     The hosts file location on Linux:
     `/etc/hosts`
4. Begin the NetBackup client installation again.

The target NetBackup server cannot be authenticated, or the server name does not match any of the host names that are listed in the server's certificate

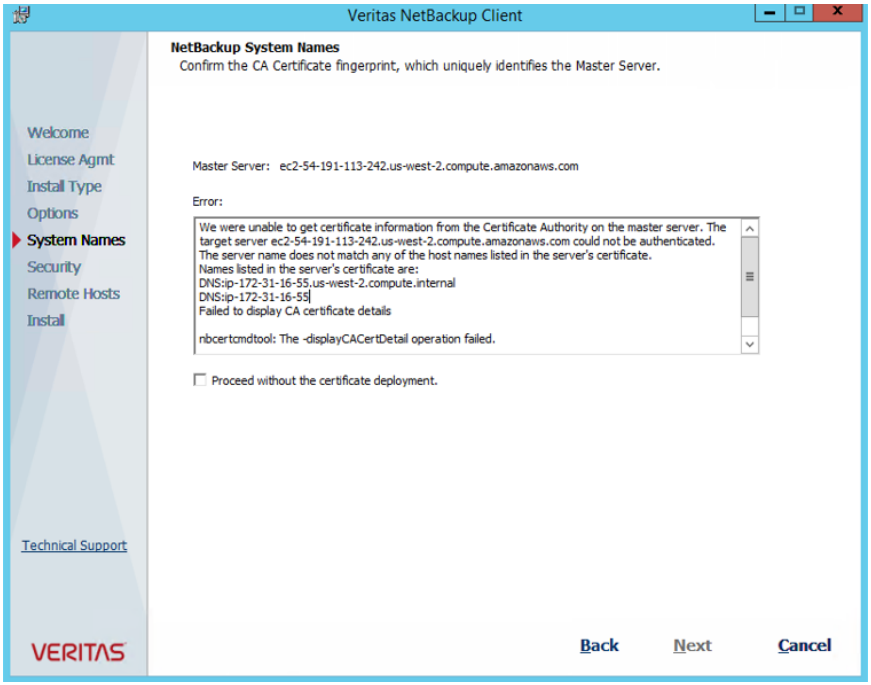

To fix this problem, do the following:

1. Click **Cancel** to cancel the installation.
2. Follow the previous steps under "Failed to display CA certificate details."
3. If the error persists, see the following article:

   Host validation fails when a NetBackup client tries to connect to the primary server
4. Begin the NetBackup client installation again.

Appendix A

Configuring services for NFS on Windows

This appendix includes the following topics:

About installing and configuring Network File System (NFS) for Granular Recovery Technology (GRT)

NetBackup Granular Recovery leverages Network File System, or NFS, to read individual objects from a database backup image. Specifically, the NetBackup client uses NFS to extract data from the backup image on the NetBackup media server. The NetBackup client uses “Client for NFS” to mount and access a mapped drive that is connected to the NetBackup media server. The NetBackup media server handles the I/O requests from the client through NBFSD.

NBFSD is the NetBackup File System (NBFS) service that runs on the media server. NBFSD makes a NetBackup backup image appear as a file system folder to the NetBackup client over a secure connection.

About configuring services for NFS on Windows 2012 or 2016 (NetBackup for VMware)

For instant recovery of virtual machines, the Services for Network File System (NFS) must be installed on Windows media servers.

Table A-1 Configuring NFS in a Windows 2012 or 2016 environment

Action	Description
Configure NFS on the media server.	On the media server do the following: ■ Stop and disable the ONC/RPC Portmapper service, if it exists. ■ Enable NFS. See "Enabling Services for Network File System (NFS) on a Windows 2012 or 2016 media server (NetBackup for VMware)" on page 361. ■ Stop the Server for NFS service. See "Disabling the Server for NFS (NetBackup for VMware)" on page 368. ■ Configure the portmap service to start automatically at server restart. Issue the following from the command prompt: `sc config portmap start= auto` This command should return the status `[SC] ChangeServiceConfig SUCCESS`.
Configure NFS on the restore host.	On the restore host, do the following: ■ Enable NFS. See "Enabling Services for Network File System (NFS) on a Windows 2012 or 2016 restore host (NetBackup for VMware)" on page 365. ■ Stop the Server for NFS service. See "Disabling the Server for NFS (NetBackup for VMware)" on page 368.

Enabling Services for Network File System (NFS) on a Windows 2012 or 2016 media server (NetBackup for VMware)

To perform VM instant recovery with a Windows 2012 or 2016 media server, you must enable Services for Network File System. When this configuration is completed, you can disable any unnecessary NFS services.

To enable Services for Network File System (NFS) on a Windows 2012 or 2016 media server

1. Open the Server Manager.
2. From the **Manage** menu, click **Add Roles and Features**.

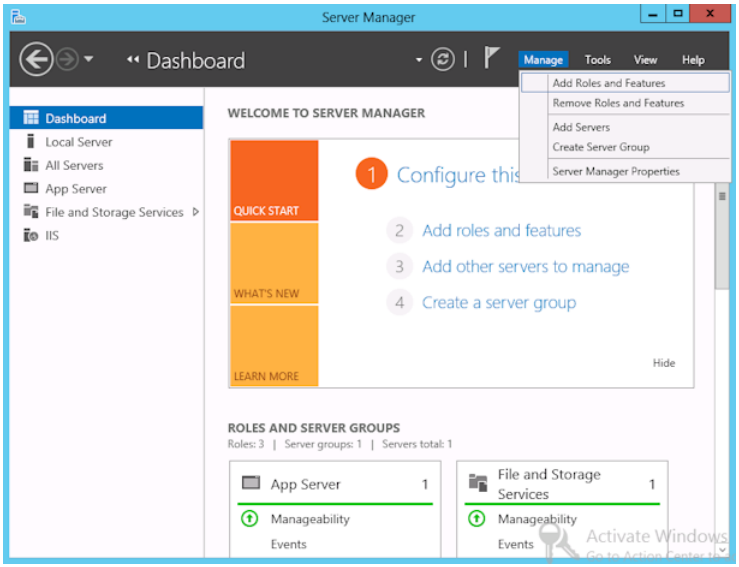

3. In the Add Roles and Features Wizard, on the **Before You Begin** page, click **Next**.

4 On the **Select installation type** page, select **Role-based or feature-based installation**.

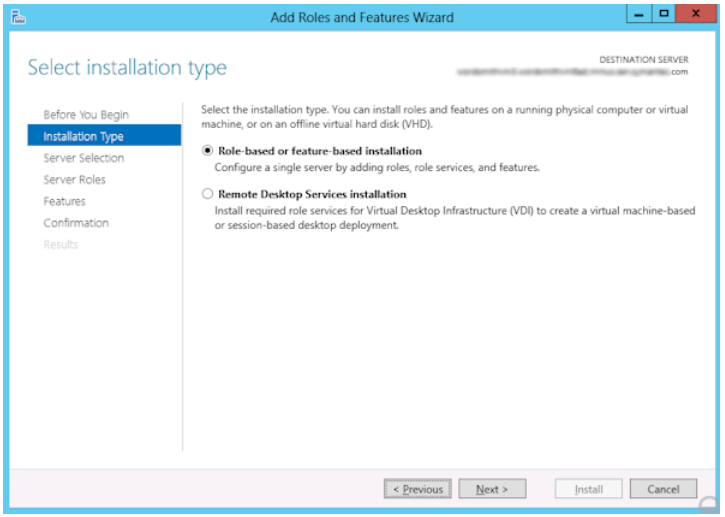

5 Click **Next**.

6 On the **Server Selection** page, click **Select a server from the server pool** and select the server. Click **Next**.

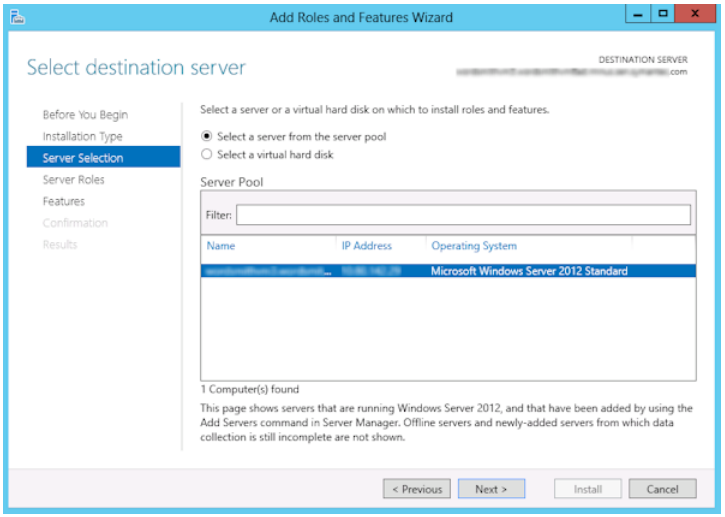

7 On the **Server Roles** page, expand **File and Storage Services** and **File and iSCSI Services**.

8 Click **File Server** and **Server for NFS**. When you are prompted, click **Add Features**. Click **Next**.

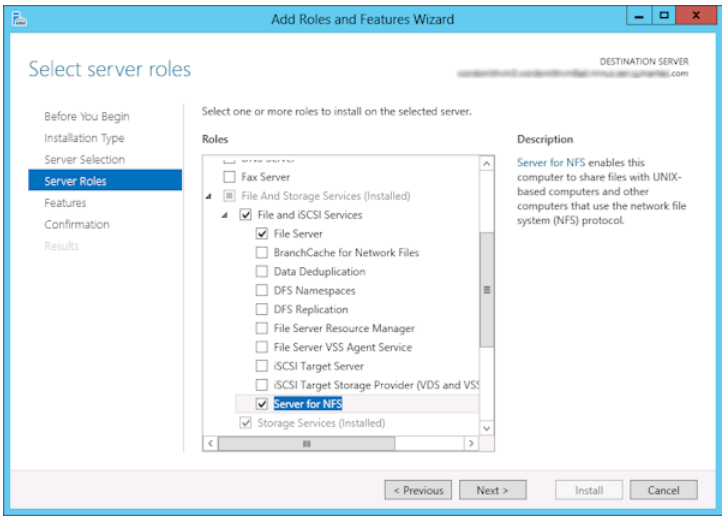

9 If the media server is also a restore host, on the **Features** page, click **Client for NFS**. Click **Next**.

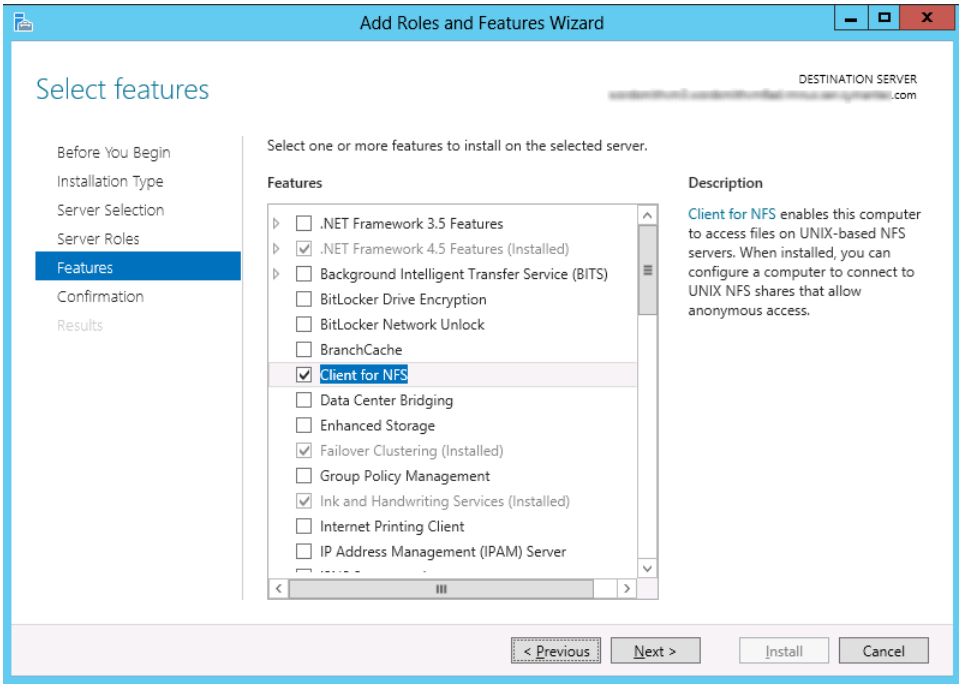

10 On the **Confirmation** page, click **Install**.

11 Disable any unnecessary services, as follows:

- If you have a single host that functions as both the media server and the restore host, you can disable the Server for NFS.
- For a host that is only the NetBackup media server, you can disable the Server for NFS and the Client for NFS.

12 Make sure that the portmap service is started and that its startup mode is set to auto.

Enabling Services for Network File System (NFS) on a Windows 2012 or 2016 restore host (NetBackup for VMware)

To perform VM instant recovery with a Windows restore host, you must enable Services for Network File System. When this configuration is complete, you can disable any unnecessary NFS services.

To enable Services for Network File System (NFS) on a Windows 2012 or 2016 restore host

1. Open the Server Manager.
2. From the **Manage** menu, click **Add Roles and Features**.

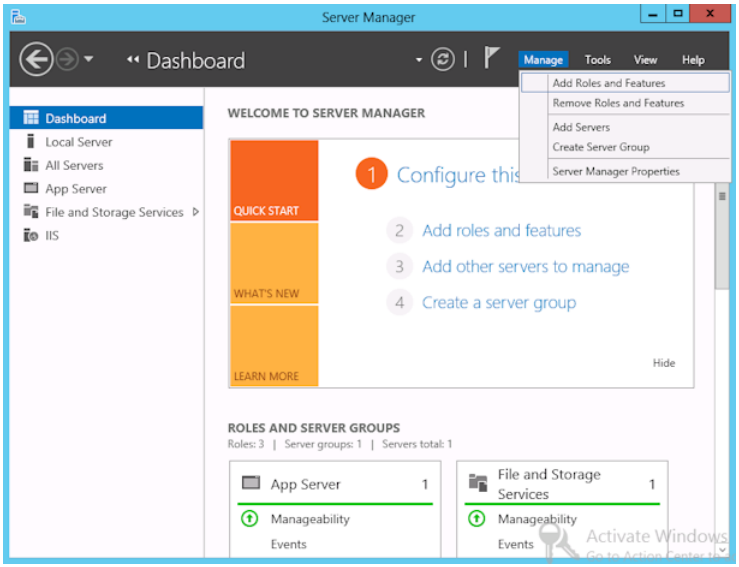

3. In the Add Roles and Features Wizard, on the **Before You Begin** page, click **Next**.

4 On the **Select installation type** page, select **Role-based or feature-based installation**.

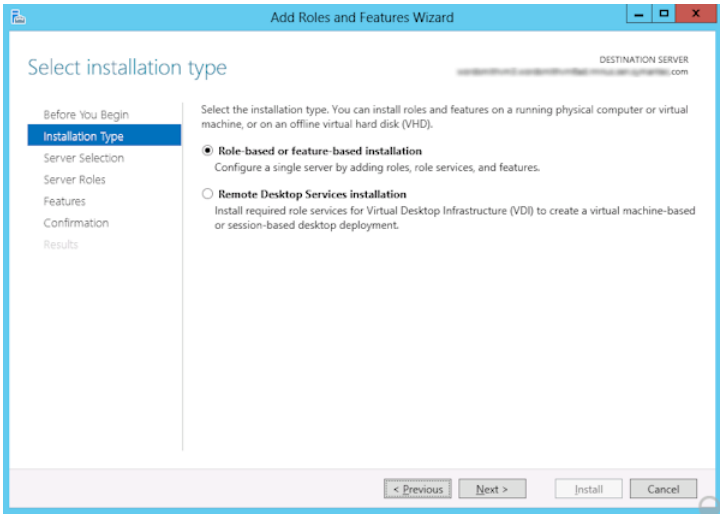

5 Click **Next**.

6 On the **Server Selection** page, click **Select a server from the server pool** and select the server. Click **Next**.

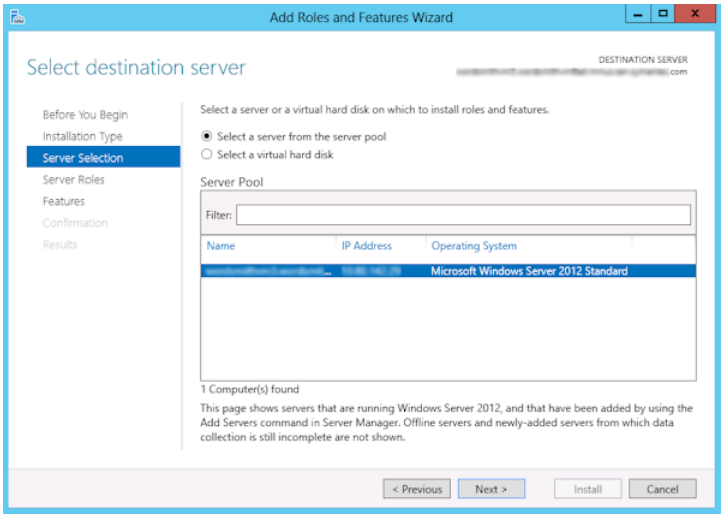

7 On the **Server Roles** page, click **Next**.

8 On the **Features** page, click **Client for NFS**. Click **Next**.

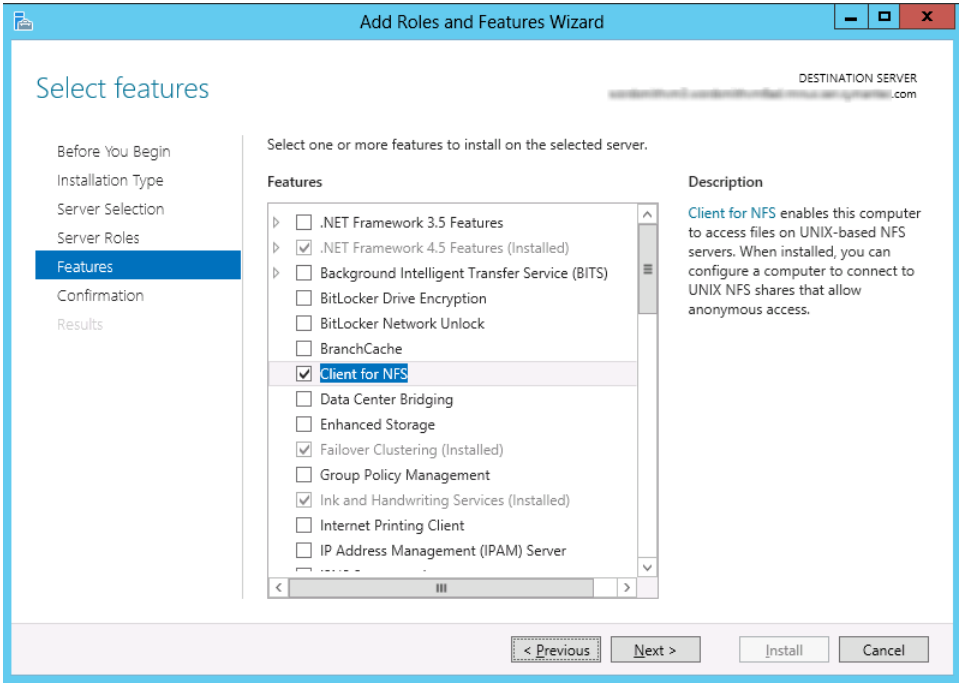

9 On the **Confirmation** page, click **Install**.

Disabling the Server for NFS (NetBackup for VMware)

To disable the Server for NFS

1 Open the Server Manager.

2 In the left pane, expand **Configuration**.

3 Click **Services**.

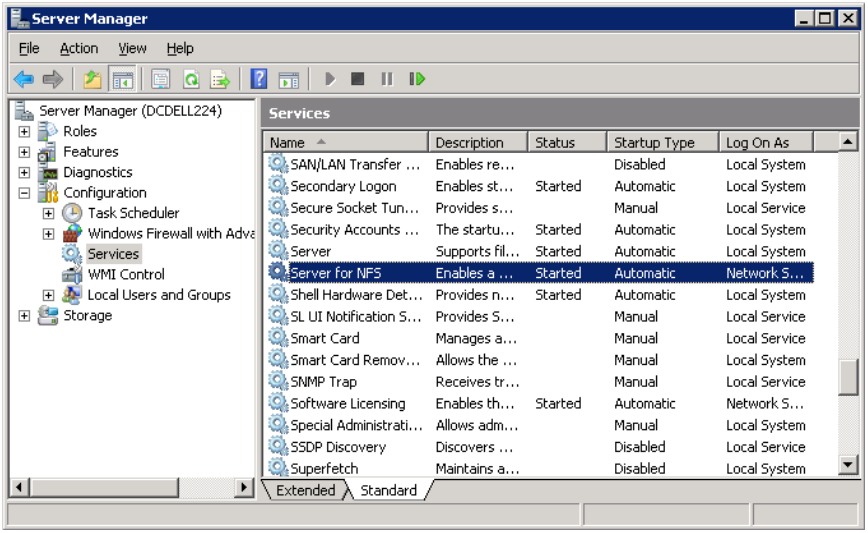

4 In the right pane, right-click on **Server for NFS** and click **Stop**.

5 In the right pane, right-click on **Server for NFS** and click **Properties**.

6 From the **Startup type** list in the **Server for NFS Properties** dialog box, click **Disabled**.

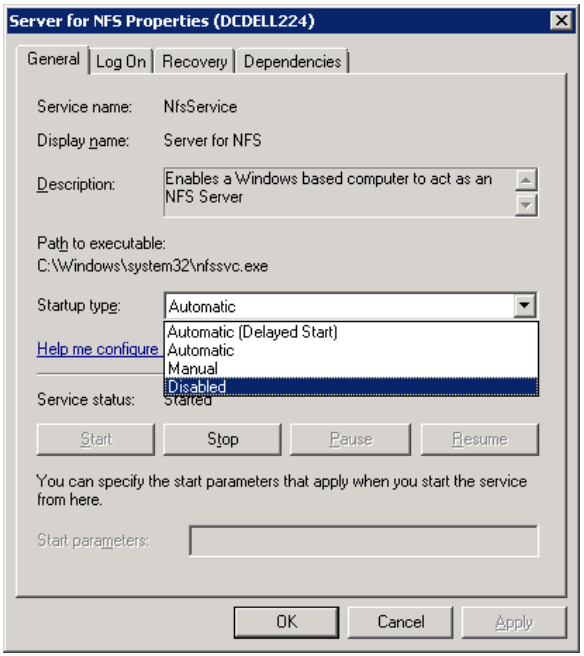

7 Click **OK**.

8 Do this procedure for each media server and for the restore host.

Disabling the Client for NFS on the media server (NetBackup for VMware)

After you enable Services for Network File System (NFS) on a host that is only a NetBackup media server, disable the Client for NFS.

To disable the Client for NFS on the NetBackup media server

1 Open the Server Manager.

2 In the left pane, expand **Configuration**.

3 Click **Services**.

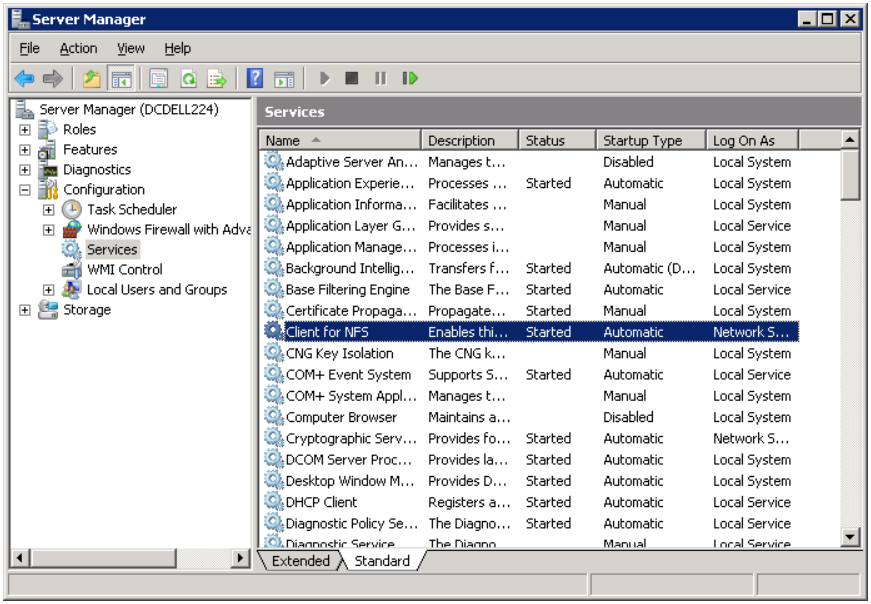

4 In the right pane, right-click on **Client for NFS** and click **Stop**.

5 In the right pane, right-click on **Client for NFS** and click **Properties**.

6 From the **Startup type** list in the **Client for NFS Properties** dialog box, click **Disabled**.

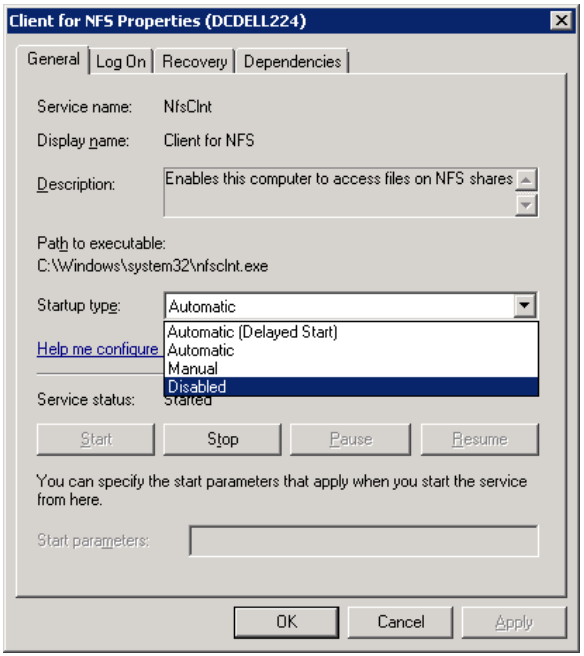

7 Click **OK**.

Configuring a UNIX media server and Windows backup or restore host for Granular Recovery Technology (NetBackup for VMware)

For backups and restores that use Granular Recovery Technology (GRT), perform the following configuration if you use a UNIX media server and Windows restore host:

- Confirm that your media server is installed on a platform that supports granular recovery.
  For more information about supported platforms, see the *NetBackup Enterprise Server and Server - OS Software Compatibility List* at the following URL:
  No other configuration is required for the UNIX media server.
- Enable or install NFS on the restore host.

See "Enabling Services for Network File System (NFS) on a Windows 2012 or 2016 restore host (NetBackup for VMware)" on page 365.

- You can configure a different network port for NBFSD.

Configuring a different network port for NBFSD (NetBackup for VMware)

`NBFSD` runs on port 7394. If another service uses the standard `NBFSD` port in your organization, you can configure the service on another port. The following procedures describe how to configure a NetBackup server to use a network port other than the default.

To configure a different network port for NBFSD (Windows server)

1. Log on as administrator on the computer where NetBackup server is installed.
2. Open Regedit.
3. Open the following key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Veritas\NetBackup\CurrentVersion\Config
```

4. Create a new DWORD value named **FSE_PORT**.
5. Right-click on the new value and click **Modify**.
6. In the **Value data** box, provide a port number between 1 and 65535.
7. Click **OK**.

To configure a different network port for NBFSD (UNIX server)

1. Log on as root on the computer where NetBackup server is installed.
2. Open the `bp.conf` file.
3. Add the following entry, where *XXXX* is an integer and is a port number between 1 and 65535.

   FSE_PORT = *XXXX*

See "Configuring a UNIX media server and Windows backup or restore host for Granular Recovery Technology (NetBackup for VMware)" on page 372.

Appendix B

Backups of VMware raw devices (RDM)

This appendix includes the following topics:

About VMware raw device mapping (RDM)

VMware raw device mapping mode (RDM) allows a virtual machine to directly access physical disks. With raw device mapping, a VMware virtual machine can use large storage devices such as disk arrays. Access to the data on an RDM disk is faster than to a fully virtualized disk (vmdk file). An RDM disk can be locally attached to the ESX server or configured on a Fibre Channel SAN.

NetBackup supports the disk arrays that are configured on a virtual machine as RDMs.

Note: NetBackup cannot back up the RDM by means of a VMware backup host.

For notes and restrictions on NetBackup support for VMware RDM, see the following Cohesity tech note:

Support for NetBackup in virtual environments

http://www.veritas.com/docs/000006177

Configurations for backing up RDMs

You can use either of the following NetBackup configurations to back up disk arrays as RDMs:

- Without Snapshot Client: Install a NetBackup client on the virtual machine. You can configure NetBackup to back up the virtual machine and any RDMs as if the client was installed on a physical host. Without Snapshot Client software on the virtual machine, the features of Snapshot Client are not available. (This configuration is not discussed in this NetBackup for VMware guide.)
- With Snapshot Client: Install a NetBackup client and Snapshot Client software on the virtual machine. Configure an alternate client backup.

About alternate client backup of RDMs

Alternate client backup of an RDM consists of the following:

- The RDM disk array contains the data to be backed up. Another host containing NetBackup client software and Snapshot Client software must have access to the disk array. This host is the alternate client. In this configuration, the virtual machine is called the primary client.
- A snapshot of the data is created on the disk array and is mounted on the alternate client. The alternate client creates a backup image from the snapshot, using original path names, and streams the image to the NetBackup media server.
- The alternate client handles the backup I/O processing; the backup has little or no effect on the virtual machine. The media server reads the snapshot data from the alternate client and writes the data to storage.
- The virtual machine and alternate client must be running the same operating system, volume manager, and file system. For each of these I/O system components, the alternate client must be at the same level as the primary client, or higher level.

Requirements for alternate client backup of RDMs

To use NetBackup Snapshot Client to back up an RDM, note the following:

- RDM devices must be configured in physical compatibility mode. You select this mode when you create the RDM. Physical compatibility mode is not configured in NetBackup.
  For an introduction to RDM, refer to your VMware documentation. For example, see the following VMware document:
  ESX Server 3 Configuration Guide
- NetBackup may require certain OS and array configuration, depending on the guest OS and the array.
- NetBackup client software must be installed on the virtual machine.
- The requirements for the NetBackup for VMware feature (a backup host and the VMware snapshot method) do not apply to backups of RDM disk arrays. To back up RDM disk arrays, you must configure a Snapshot Client alternate client backup.

Configure an alternate client backup of RDMs

This procedure highlights key points in creating a NetBackup alternate client backup of a disk array that is configured as an RDM.

To create an alternate client policy for a disk array that is configured as an RDM

1. Open the NetBackup web UI.
2. Select a policy type that is appropriate for the OS of the virtual machine and for the type of data to back up.
3. On the policy **Attributes** tab, click **Perform snapshot backups** and **Perform off-host backup**.
4. Select **Alternate client** from the **Use** list. Do not select **VMware backup host**.

   In the **Machine** field, enter the name of the host that is configured as an off-host backup computer (the alternate client).
5. Click **Snapshot options**.
6. Select a snapshot method.

   The VMware method does not apply to alternate client backup and is not available in the list.

   Select a snapshot method that is appropriate for the volume or array. For example:

   - The HP_EVA_Snapclone method or other EVA method for an HP EVA array.

- The EMC_CLARiiON_Snapview_Clone or other CLARiiON method for an EMC CLARiiON array.
- FlashSnap.
  For FlashSnap, the following must be installed: VxVM 3.2 or later for UNIX, VxVM 4.0 or later for Linux and AIX, or VxVM 3.1 or later for Windows. Also, volumes must be configured over the primary host's disks. The VxVM FlashSnap license must also be installed.
- VSS (for Windows guest operating systems only).

The array may require additional OS and NetBackup configuration.

7 If required by an array snapshot method that you selected in the previous step, specify the **Snapshot Resources**.

8 In the policy's **Clients** list, select the virtual machine on which the array is configured as an RDM.

9 In the policy's **Backup selections** tab, specify the disk that you want to back up, or the files or volumes that reside on the disk.