

NetBackup™ Web UI Cloud Administrator's Guide

Release 9.0

VERITAS™

NetBackup Web UI Cloud Administrator's Guide

Last updated: 2021-01-05

Legal Notice

Copyright © 2021 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies LLC and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. Veritas Technologies LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
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 Veritas 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 Veritas 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 Veritas community site:

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

Veritas Services and Operations Readiness Tools (SORT)

Veritas 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	Introducing the NetBackup web user interface	6
	6
	About the NetBackup web UI	6
	Terminology	8
	Sign in to the NetBackup web UI	10
	Sign out of the NetBackup web UI	12
Chapter 2	Monitoring and notifications	13
	13
	The NetBackup dashboard	13
	Monitoring jobs	14
	Filter jobs in the job list	14
Chapter 3	Managing and protecting cloud assets	16
	16
	About protecting cloud assets	16
	Limitations and considerations	18
	AWS and Azure government cloud support	19
	About protecting Microsoft Azure resources using resource groups	
	20
	Before you begin	21
	Limitations and considerations	21
	About resource group configurations and outcome	21
	Troubleshoot resource group permissions	24
	CLOUD_AUTODISCOVERY_INTERVAL option for NetBackup servers	
	25
	Configure snapshot replication	25
	Protect applications in-cloud with application consistent snapshots	
	27
	Configure CloudPoint servers in NetBackup	29
	Configure a third-party CA certificate	30
	Add a CloudPoint server	31
	Add a cloud provider for a CloudPoint server	32
	Associate media servers with a CloudPoint server	35
	Discover assets on CloudPoint server	35
	Edit a CloudPoint server	36

	Enable or disable a CloudPoint server	36
Chapter 4	Recovering cloud assets	37
	Recover a cloud asset to its original location	37
	Recover a cloud asset to an alternate location	38
	Perform rollback recovery of cloud assets	38
Chapter 5	Troubleshooting protection and recovery of cloud assets	40
	Troubleshoot cloud workload protection issues	40
Chapter 6	Performing granular restore	45
	About granular restore	45
	Supported environment list	46
	List of supported file systems	47
	Before you begin	47
	Limitations and considerations	49
	Restoring files and folders cloud virtual machines	50
	Restoring volumes on cloud virtual machines	52
	Troubleshooting snapshot restore process for Microsoft Azure-specific cloud	53

Introducing the NetBackup web user interface

This chapter includes the following topics:

- [About the NetBackup web UI](#)
- [Terminology](#)
- [Sign in to the NetBackup web UI](#)
- [Sign out of the NetBackup web UI](#)

About the NetBackup web UI

The NetBackup web user interface provides the following features:

- Ability to access the master server from a web browser, including Chrome and Firefox. For details on supported browsers for the web UI, see the [NetBackup Software Compatibility List](#).
Note that the NetBackup web UI may behave differently for different browsers. Some functionality, for example a date picker, may not be available on all browsers. These inconsistencies are due to the capabilities of the browser and not because of a limitation with NetBackup.
- A dashboard that displays a quick overview of the information that is important to you.
- Role-based access control (RBAC) that lets the administrator configure user access to NetBackup and to delegate the tasks such as security, storage management, or workload protection.
- Management of NetBackup security settings, certificates, API keys, and user sessions.

- Protection of assets is achieved through protection plans, job management, and visibility of the protection status of assets. Alternatively, policy management is also available for a limited number of policy types.
- Workload administrators can create protection plans, subscribe assets to the protection plans that meet the SLO, monitor protection status, and perform self-service recovery of virtual machines. The web UI supports the following workloads:
 - Cloud
 - Microsoft SQL Server
 - Oracle
 - Red Hat Virtualization (RHV)
 - VMware
- Usage reporting tracks the size of backup data on your master servers. You can also easily connect to Veritas NetInsights Console to view and manage NetBackup licensing.

Note: The NetBackup web UI is best viewed at a 1280x1024 or higher screen resolution.

Access control in the NetBackup web UI

NetBackup uses role-based access control to grant access to the web UI. Access control is accomplished through roles.

- A role defines the operations that a user can perform and the access that the user has to any workload assets, protection plans, or credentials. A user can have multiple roles, allowing for full and for flexible customization of user access.
- RBAC is only available for the web UI and the APIs.
Other access control methods for NetBackup are not supported for the web UI and APIs, with the exception of Enhanced Auditing (EA).

Monitor NetBackup jobs and events

The NetBackup web UI lets administrators more easily monitor NetBackup operations and events and identify any issues that need attention.

- The dashboard displays an overview of NetBackup jobs, certificates, tokens, security events, and usage reporting.
The dashboard widgets that display depend on a user's RBAC role and permissions.

- Email notifications can be configured so administrators receive notifications when job failures occur. NetBackup supports any ticketing system that can receive inbound email.

Protection plans: One place to configure schedules, storage, and storage options

Protection plans offer the following benefits:

- In addition to schedules for backups, a protection plan can also include a schedule for replication and long-term retention.
- When you select from your available storage, you can see any additional features available for that storage.
- A default workload administrator can create and manage protection plans, including the backup window and retention.
See *NetBackup Web UI Administrator's Guide* for details on the roles permissions.
- A default workload administrator can select the protection plans to use to protect assets or intelligent groups.

Self-service recovery

The NetBackup web UI makes it easy for a workload administrator to recover VMs or databases. For the workloads that support the instant access feature, users can mount a snapshot for immediate access to a VM's files or to a database.

Terminology

The following table describes the concepts and terms that are introduced with the new web user interface.

Table 1-1 Web user interface terminology and concepts

Term	Definition
Administrator	A user that has complete access and permissions to NetBackup and all of the interfaces, including the NetBackup web UI. The root, administrator, and Enhanced Auditing user all have complete access to NetBackup. In the <i>NetBackup Web UI</i> guides, the term <i>NetBackup administrator</i> also refers to a user that has full permissions for NetBackup. Usually in reference to a user of the NetBackup Administration Console. Also see <i>role</i> .
Asset group	See <i>intelligent group</i> .

Table 1-1 Web user interface terminology and concepts (*continued*)

Term	Definition
Asset	The data to be protected, such as physical clients, virtual machines, and database applications.
Backup now	An immediate backup of an asset. NetBackup performs a one-time, full backup of an asset using the selected protection plan. This backup does not affect any scheduled backups.
Classic policy	In the NetBackup web UI, indicates that a legacy policy protects the asset. Legacy policies are created with the NetBackup Administration Console.
External certificate	A security certificate that is issued from any CA other than NetBackup.
Intelligent group	<p>Allows NetBackup to automatically select assets for protection based on the criteria (queries) that you specify. An intelligent group automatically stays up-to-date with changes in the production environment. These groups are also referred to as asset groups.</p> <p>For VMware and RHV, these groups appear under the tab Intelligent VM groups.</p>
Instant access	An instant access VM or database that is created from a NetBackup backup image is available almost instantaneously, achieving a near-zero recovery time objective. NetBackup mounts the snapshot directly on the backup storage device and the snapshot is treated as a normal VM or database.
NetBackup certificate	A security certificate that is issued from the NetBackup CA.
Protection plan	A protection plan defines when backups are performed, how long the backups are retained, and the type of storage to use. Once a protection plan is set up, assets can be subscribed to the protection plan.
RBAC	<p>Role-based access control. Administrators can delegate or limit access to the NetBackup web UI through the roles that are configured in RBAC.</p> <p>Note: The roles that you configure in RBAC do not control access to the NetBackup Administration Console or the CLIs.</p>

Table 1-1 Web user interface terminology and concepts (*continued*)

Term	Definition
Role	For RBAC, defines the operations that a user can perform and the assets or objects that they can access. For example, you can configure a role to manage recovery of specific databases and the credentials that are needed for backups and restores.
Storage	The storage to which the data is backed up, replicated, or duplicated (for long-term retention).
Subscribe, to a protection plan	The action of selecting an asset or an asset group to subscribe to a protection plan. The asset is then protected according to the schedule in the plan. The web UI also refers to <i>Subscribe</i> as <i>Add protection</i> .
Unsubscribe, from a protection plan	<i>Unsubscribe</i> refers to the action of removing protection or removing an asset or asset group from a plan.
Workload	The type of asset. For example, VMware, RHV, or Cloud.

Sign in to the NetBackup web UI

Authorized users can sign in to a NetBackup master server from a web browser, using the NetBackup web UI. The following sign-in options are available:

- [Sign in with a username and password](#)
- [Sign in with a certificate or smart card](#)
- [Sign in with single sign-on \(SSO\)](#)

Sign in with a username and password

Only authorized users can sign in to NetBackup web UI. Contact your NetBackup security administrator for more information.

To sign in to a NetBackup master server using a username and password

- 1 Open a web browser and go to the following URL.

`https://masterserver/webui/login`

The *masterserver* is the host name or IP address of the NetBackup master server that you want to sign in to.

- 2 Enter your credentials and click **Sign in**.

For example:

For this type of user	Use this format	Example
Local user	<i>username</i>	jane_doe
Windows user	<i>DOMAINusername</i>	WINDOWS\jane_doe
UNIX user	<i>username@domain</i>	john_doe@unix

Sign in with a certificate or smart card

You can sign in to NetBackup web UI with a smart card or digital certificate if you are an authorized user. Contact your NetBackup security administrator for more information.

To use a digital certificate that is not on a smart card, you must first upload the certificate to the browser's certificate manager. See the browser documentation for instructions or contact your certificate administrator for more information.

To sign in with a certificate or smart card

- 1 Open a web browser and go to the following URL.

`https://masterserver/webui/login`

The *masterserver* is the host name or IP address of the NetBackup master server that you want to sign in to.

- 2 Click **Sign in with certificate or smart card**.
- 3 When your browser prompts you, select the certificate.

Sign in with single sign-on (SSO)

You can sign in to NetBackup web UI with the single sign-on (SSO) option if SAML is configured as an identity provider in your NetBackup environment. Contact your NetBackup security administrator for more information.

To sign in to a NetBackup master server using SSO

- 1** Open a web browser and go to the following URL.

`https://masterserver/webui/login`

The *masterserver* is the host name or IP address of the NetBackup master server that you want to sign in to.

- 2** Click **Sign in with single sign-on**.
- 3** Follow the steps as provided by your administrator.

On subsequent logons, NetBackup signs you in automatically to the master server.

Sign out of the NetBackup web UI

Note that NetBackup automatically signs you out of the web UI after 24 hours, which is the maximum time that is allowed for a user session. After that time NetBackup requires that you sign in again. You can also sign out if you want to change the sign-in option that you want to use (username and password, smart card, or single sign-on (SSO)).

To sign out of the NetBackup web UI

- ◆ On the top right, click the profile icon and click **Sign out**.

Monitoring and notifications

This chapter includes the following topics:

- [The NetBackup dashboard](#)
- [Monitoring jobs](#)
- [Filter jobs in the job list](#)

The NetBackup dashboard

The NetBackup dashboard provides a quick view of the details that are related to your role in your organization.

Table 2-1 The NetBackup dashboard

Dashboard widget	Description
Jobs	Lists job information, including the number of active and queued jobs and the status of attempted and completed jobs.
Certificates	<p>Displays the information about the NetBackup host ID-based security certificates or the external certificates in your environment.</p> <p>For external certificates, the following information is shown for NetBackup 8.2 and later hosts:</p> <ul style="list-style-type: none">▪ Total hosts. The total number hosts. The hosts must be online and able to communicate with NetBackup master server.▪ Missing. The number hosts that do not have an external certificate enrolled.▪ Valid. The number of hosts that have an external certificate enrolled.▪ Expired. The number of hosts with expired external certificates.

Table 2-1 The NetBackup dashboard (*continued*)

Dashboard widget	Description
Tokens	Displays the information about the authorization tokens in your environment.
Security events	The Access history view includes a record of logon events. The Audit events view includes the events that users initiate on the NetBackup master server.
Usage reporting	Lists the size of the backup data for the NetBackup master servers in your organization. This reporting is useful to track capacity licensing. Use the drop-down lists in the top right to select the time period and the view that you want to display. Click on a server name to see specific details for that server.

Monitoring jobs

Use the **Jobs** node to monitor the jobs in your NetBackup environment and view the details for a specific job.

To monitor a job

- 1 On the left, click **Activity monitor > Jobs**.
- 2 Click on a job name that you want to view.
On the **Overview** tab you can view information about a job.
 - The **File List** contains the files that are included in the backup image.
 - The **Status** section shows the status and the status codes that are related to the job. Click the status code number to view information about this status code in the Veritas Knowledge Base.
See the [NetBackup Status Codes Reference Guide](#).
- 3 Click the **Details** tab to view the logged details about a job. You can filter the logs by error type using the drop-down menu.
See [“Filter jobs in the job list”](#) on page 14.

Filter jobs in the job list

You can filter the jobs to display the jobs in a specific state. For example, you can display all of the active jobs or all of the suspended jobs.

To filter the job list

- 1 Click **Jobs**.
- 2 Above the job list, click the **Filter** option.

- 3 In the **Filter** window, select a filter option to dynamically change the jobs that are displayed. The filter options are as follows:
 - **All**
 - **Active**
 - **Done**
 - **Failed**
 - **Incomplete**
 - **Partially Successful**
 - **Queued**
 - **Successful**
 - **Suspended**
 - **Waiting for Retry**
- 4 Click **Apply Filters**.
- 5 To remove the selected filters, click **Clear All**.

Managing and protecting cloud assets

This chapter includes the following topics:

- [About protecting cloud assets](#)
- [Limitations and considerations](#)
- [AWS and Azure government cloud support](#)
- [About protecting Microsoft Azure resources using resource groups](#)
- [CLOUD_AUTODISCOVERY_INTERVAL option for NetBackup servers](#)
- [Configure snapshot replication](#)
- [Protect applications in-cloud with application consistent snapshots](#)
- [Configure CloudPoint servers in NetBackup](#)

About protecting cloud assets

Using NetBackup, you can now protect your in-cloud workloads. The cloud data protection framework leverages the CloudPoint infrastructure to drive faster proliferation of cloud providers. Starting with 9.0, CloudPoint can now protect assets that use IPv6 as network communication channel. IPv6 is supported only in AWS commercial and Gov Cloud. It is not supported for Azure commercial Cloud, Azure Gov Cloud, and GCP.

The following table describes the tasks.

Table 3-1 Configuring protection for cloud assets

Task	Description
<p>Before you begin ensure that you have the appropriate permission.</p>	<p>To manage and protect cloud assets in the web UI you must have the workload administrator role or similar permissions. Contact the NetBackup security administrator.</p> <p>See the NetBackup Web UI Administrator's Guide.</p> <p>Note: For managing hosted applications, you need Manage Assets and Manage Protection Plans permissions.</p>
<p>Deploy CloudPoint</p>	<p>Install CloudPoint in your environment.</p> <p>See "Add a CloudPoint server" on page 31.</p> <p>Review CloudPoint and NetBackup limitations.</p> <p>See "Limitations and considerations" on page 18.</p>
<p>Configure the CloudPoint server using the NetBackup Administration Console</p>	<p>Register the CloudPoint server in NetBackup.</p> <p>See, <i>Veritas NetBackup Snapshot Client Administrator's Guide</i>.</p>
<p>Add a configuration</p>	<p>All the supported cloud providers are displayed in the web UI.</p> <p>You need to add the cloud account (configure the cloud plug-in) for the cloud provider you need. You can create multiple configurations for each provider.</p> <p>See "Add a cloud provider for a CloudPoint server" on page 32.</p> <p>For Amazon, you can choose to use IAM role.</p> <p>See "IAM Role for AWS Configuration" on page 34.</p>
<p>Asset discovery</p>	<p>NetBackup retrieves the cloud assets pertaining to the cloud accounts that are configured in NetBackup. Assets are populated in NetBackup asset DB.</p> <p>By default, asset discovery happens every 2 hours and is configurable.</p> <p>In case of applications, you can set discovery interval between 15 minutes to 45 minutes.</p> <p>See "CLOUD_AUTODISCOVERY_INTERVAL option for NetBackup servers" on page 25.</p>

Table 3-1 Configuring protection for cloud assets (*continued*)

Task	Description
Create a snapshot only protection plan	<p>Create a snapshot only protection plan. A protection plan is used to schedule backup start windows.</p> <p>See the NetBackup Web UI Administrator's Guide.</p> <p>You can also configure the protection plan for snapshot replication. See "Configure snapshot replication" on page 25.</p>
Choose to protect a virtual machine, application, or volume	<p>For each cloud provider, a list of discovered assets is displayed. Add the assets to a protection plan.</p> <p>See the NetBackup Web UI Administrator's Guide.</p> <p>You can also choose to protect application using application consistent snapshots. See "Protect applications in-cloud with application consistent snapshots" on page 27.</p>
Recover cloud assets	<ul style="list-style-type: none"> ■ You can recover the assets using the recovery points. See "Recover a cloud asset to its original location" on page 37. See "Recover a cloud asset to an alternate location" on page 38. See "Perform rollback recovery of cloud assets " on page 38. ■ You can also restore the assets using the <code>nbcloudrestore</code> CLI utility. <ul style="list-style-type: none"> Note: Do not use the <code>bprestore</code> CLI for restores <p>See the NetBackup Commands Reference Guide.</p>
Troubleshooting	<p>See "Troubleshoot cloud workload protection issues" on page 40.</p>

Limitations and considerations

Consider the following when protecting cloud workloads

- Deletion of CloudPoint host entry and its associated plug-ins is not supported in NetBackup.

If you delete plug-ins that are configured in NetBackup, you cannot recover any CloudPoint images that are associated with that plug-in.
- Review the *Veritas CloudPoint Install and Upgrade Guide* for information on the capabilities of CloudPoint.

- NetBackup integration is not supported with the CloudPoint freemium version.
- If you have a previous installation of CloudPoint, Veritas recommends that you upgrade the CloudPoint server and not reinstall it.
 If you do reinstall the CloudPoint server, you need to reconfigure the CloudPoint server and perform all the protection-related steps.
- When you configure a CloudPoint server using port 0, the default value is used.
- After CloudPoint server is added, the host machine tries to use the IPv6 address to discover assets on cloud. If the IPV6 address is found on the host, the application is configured to use it. If an IPv6 address is not found, the IPv4 address is used.
- Cloudpoint 8.3 only supports IPv6 when it is installed on Ubuntu 18.04 or later and RHEL Operating Systems. If you are using Ubuntu 16.04, you must first upgrade the OS to Ubuntu 18.04 to use IPv6.
- For CloudPoint server, enhanced auditing is not supported. Thus, when you add or update a CloudPoint server, with non-root but NetBackupAdmin rights, during auditing the user is shown as root.
- If you deploy CloudPoint using the CloudFormation template, when you register the on-host with the CloudPoint node using the command, the IP address used must be private IP and not public IP.

AWS and Azure government cloud support

Starting with 8.3, the CloudPoint server can discover Amazon Web Services and Microsoft Azure US Government cloud workloads. After the CloudPoint server is added to NetBackup, you can protect the workloads by NetBackup. NetBackup is compliant with the regulatory requirements including IPv6 support to deploy CloudPoint on the AWS and Azure US government cloud workloads.

After you configure AWS or Azure US Government cloud, the AWS and Azure agent service is created which discovers the cloud assets based on provided region. The discovered assets are displayed in NetBackup. Currently, only workloads from selected regions and mapped endpoint are discovered and protected. For the same CloudPoint host, you cannot use a combination of public and government clouds.

An error might occur if you update a cloud plug-in when the plug-in assets operations are in-progress.

CloudPoint supports the following GovCloud (US) regions:

Cloud provider

Amazon Web Services

GovCloud (US) regions

- us-gov-east-1
- us-gov-west-1

Microsoft Azure

- US Gov Arizona
- US Gov Texas
- US Gov Virginia

For information about configuring AWS and Microsoft Azure, See [“Add a cloud provider for a CloudPoint server”](#) on page 32.

About protecting Microsoft Azure resources using resource groups

NetBackup lets you define a peer Resource Groups snapshot destination for every resource group that contains protected virtual machines and volumes.

All resources in Microsoft Azure are associated to a resource group. After a snapshot is created, it is associated to a resource group. Also, each resource group is associated to a region. See,

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/manage-resource-groups-portal>

CloudPoint creates a snapshot and places the snapshot in resource group to which the resource belongs even under the following conditions:

- If you don't provide a prefix for a resource group
- Peer resource groups are not created
- You allow the snapshots to get created

You can configure the settings to place the snapshots in different resource group than the resource group that is associated with the resource. However, note the following important points:

- The peer resource group must be in the same region as the region of the resource group of the resource.
- If a peer resource group is not found, the configurations determine whether the snapshots creation succeeds or fails.

To enable this feature, you must create peer resource groups. CloudPoint then appends the prefix of the resource group that is associated with the resource. When a snapshot is created, the peer resource group name is derived based on the prefix and the resource group to which the resource is associated.

Before you begin

- The peer resource groups must be available for resources that are being protected using the resource group.
- Regions of a plugin configuration must not overlap with another configuration if a prefix is specified.

Limitations and considerations

- Only alphanumeric characters, periods, underscores, hyphen, or parenthesis are allowed in the resource group names.
- The prefix length must be less than 89 characters.
- You cannot use characters that Azure configuration does not allow for resource group naming conventions.

About resource group configurations and outcome

The following table lists scenarios for virtual machines and resource group setup, resource configuration, and outcome.

Table 3-2 Configurations and outcome

Resource group prefix	Protect assets even if prefixed Resource Groups are not found check box	Outcome
Not specified	Not selected	NetBackup associates the newly created snapshots in resource group of the resource.
Specified	Not selected	NetBackup creates new the snapshots and associates the snapshots to the peer resource group if the following conditions are met: <ul style="list-style-type: none"> ■ The peer resource group is created. ■ The peer resource group is in the same region as the resource group. If the conditions are not met, snapshot jobs fail.

Table 3-2 Configurations and outcome (*continued*)

Resource group prefix	Protect assets even if prefixed Resource Groups are not found check box	Outcome
Specified	Selected	<p>NetBackup creates new snapshots and associates the snapshots to the peer resource group if the following conditions are met:</p> <ul style="list-style-type: none"> ■ The peer resource group is created. ■ The peer resource group is in the same region as the resource group. <p>If a peer resource group is not created or is in a different region then the newly created snapshot is associated to the resource group of the resource that is protected.</p>

Examples of resource group configurations

The following table lists the examples for resource group configurations.

Table 3-3 Example configurations

Conditions	Configurations	Result
<ul style="list-style-type: none"> ■ OS and all disks are in the same resource group. ■ Peer resource group is named correctly. ■ Peer resource is located in the same region as resource group of resource. 	<ul style="list-style-type: none"> ■ Resource Group Prefix value is provided. ■ The Protect assets even if prefixed Resource Groups are not found check box is selected. 	Snapshots are created in the peer resource group.

Table 3-3 Example configurations (*continued*)

Conditions	Configurations	Result
<ul style="list-style-type: none"> ■ OS and all disks are in separate resource groups. ■ Peer resource groups are named correctly. ■ Peer resources are located in the same region as resource groups of resources. 	<ul style="list-style-type: none"> ■ Resource Group Prefix value is provided. ■ The Protect assets even if prefixed Resource Groups are not found check box is selected. 	Snapshots are created in the peer resource group.
<ul style="list-style-type: none"> ■ OS and all disks are in the same resource group. ■ Peer resource group is created in a different region from the resource group of the resource. 	<ul style="list-style-type: none"> ■ Resource Group Prefix value is provided. ■ The Protect assets even if prefixed Resource Groups are not found check box is selected. 	The snapshots are created in original resource group not the peer resource group.
<ul style="list-style-type: none"> ■ OS and all disks are in the same resource group. ■ Peer resource group is not created. 	<ul style="list-style-type: none"> ■ Resource Group Prefix value is provided. ■ The Protect assets even if prefixed Resource Groups are not found check box is selected. 	The snapshots are created in original resource group not the peer resource group.
<ul style="list-style-type: none"> ■ OS and all disks are in separate resource groups, RG1 and RG2. ■ Peer resource groups RG1 is named correctly and located in the same region as the resources. ■ Peer resources group RG2 is not created. 	<ul style="list-style-type: none"> ■ Resource Group Prefix value is provided. ■ The Protect assets even if prefixed Resource Groups are not found check box is selected. 	Snapshots are created in the peer resource group of RG1 and original resource group RG2.
<ul style="list-style-type: none"> ■ OS and all disks are in same resource group. ■ Peer resource groups are named correctly. ■ Peer resources group is located different region than the resource group of resources. 	<ul style="list-style-type: none"> ■ Resource Group Prefix value is provided. ■ The Protect assets even if prefixed Resource Groups are not found check box is not selected. 	Snapshots are not created and the job fails.

Table 3-3 Example configurations (*continued*)

Conditions	Configurations	Result
<ul style="list-style-type: none"> ■ OS and all disks are in the same resource group. ■ Peer resource group is not created. 	<ul style="list-style-type: none"> ■ Resource Group Prefix value is provided. ■ The Protect assets even if prefixed Resource Groups are not found check box is not selected. 	<p>Snapshots are not created and the job fails.</p>
<ul style="list-style-type: none"> ■ OS and all disks are in separate resource groups, RG1 and RG2. ■ Peer resource groups of RG1 and RG2 that is, snapRG1 and snapRG2 are in different regions. ■ Peer resource group snapRG1 is located in the same region as the resource group RG1. ■ The peer resource group snapRG2 is located in a different region than resource group RG2. 	<ul style="list-style-type: none"> ■ Resource Group Prefix value is provided. ■ The Protect assets even if prefixed Resource Groups are not found check box is not selected. 	<p>Snapshots are not created and the job fails.</p>

Troubleshoot resource group permissions

If appropriate permissions are not assigned to the resource group, the snapshot creation fails for Azure resources that are associated to resource groups.

Workaround:

To resolve this issue, perform the following steps:

1. Navigate to <https://portal.azure.com/#blade/HubsExtension/BrowseResourceGroups>.
2. Click on the resource group, that is to be used in the snapshot.
3. Click on **Access control (IAM)**.
4. Click on **Add Role Assignment**.
5. Select **Role as Owner, Assign Access to as User**, and select the **Application (created for CloudPoint, to make API calls)**.
6. Save and try to backup again.

CLOUD_AUTODISCOVERY_INTERVAL option for NetBackup servers

This option controls how often NetBackup scans the CloudPoint servers to discover cloud assets to display in NetBackup.

Table 3-4 CLOUD_AUTODISCOVERY_INTERVAL information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>Note: These commands require administrator privilege on the NetBackup master server. For assistance, contact the NetBackup administrator.</p> <p>The default is 2 hours. The minimum is 2 hours, the maximum 1 year.</p> <p>Use the following format:</p> <pre>CLOUD_AUTODISCOVERY_INTERVAL = number of seconds</pre> <p>For example:</p> <pre>CLOUD_AUTODISCOVERY_INTERVAL = 100000</pre> <p>This entry should appear only once in the configuration file.</p> <p>Note: After changing this option, stop and restart the NetBackup services.</p>

Configure snapshot replication

You can choose to replicate snapshots cloud assets from the primary location to a remote or a secondary location. The snapshot management servers (CloudPoint) support cross-region and cross account replication. With snapshot replication you can achieve the following:


- Maintain a copy of cloud assets at a different destination for long-term retention and auditing requirements.

- Recover cloud assets from the replicated copies from another region in case there is a region outage.
- Recover cloud assets from the replicated copies from another account in case the user account is compromised.

Configuration

Review the following information to configure snapshot replication:

- You can configure snapshot replication when you create a protection plan. See the [NetBackup™ Web UI Backup Administrator's Guide](#).



- For cross account replication, you need to establish a trust relationship between the source and the target account. For more details, refer to the *Across AWS Accounts Using IAM Roles* related information in the *Amazon Web Services* documentation.

Considerations

Consider the following when you configure cloud snapshot replication:

- In a single protection plan, replication to only one destination region is supported.
- Even if multiple schedules are configured, the replication destination region that is configured is applied to all the schedules.
- Cloud snapshot replication is supported only for Amazon cloud providers.

Asset protection criteria

Consider the following before adding cloud assets to a protection plan that is configured for cloud snapshot replication:

- Assets must be added to a protection plan that replicates snapshots to a different region.
 For example, assets residing in region 'aws_account_1-us-east-1' cannot be subscribed to a protection plan replicating to the same region 'aws_account_1-us-east-1'.
- Assets can be replicated to a different account in the same region.
 For example, assets residing in region 'aws_account_1-us-east-1' can be subscribed to a protection plan replicating to the same region but different account 'aws_account_2-us-east-1'.
- Assets that are discovered by a snapshot management server must be replicated to the region that is discovered by the same snapshot management server.
 For example, assets that are discovered by snapshot management server 'CP1' cannot be subscribed to a protection plan replicating to a region that is discovered by snapshot management server 'CP2'.
- Only Amazon assets can be subscribed to a protection plan that is configured for cloud snapshot replication.

Manage concurrent snapshots replications

For better performance, you can tune the number of concurrent snapshot replications. Amazon has different limits for each asset type to do concurrent snapshot replications to a single destination region. For example, RDS has a limit for 5, EBS has a limit for 5, and EC2 has a limit for 50. For more details refer to *Copy Snapshot* related information in the *Amazon Web Services* documentation.

In NetBackup this limit is defined using the following parameter in the `bp.conf` file:

```
MAX_CLOUD_SNAPSHOT_REPLICATION_JOBS_PER_DESTINATION
```

The default value is 5.

Protect applications in-cloud with application consistent snapshots

You can take application consistent (point-in-time) snapshots of the applications that are deployed on virtual machines in cloud. This lets you perform a point-in-time recovery of applications.

You can perform original location and alternate location restores for these workloads.

For alternate location restore, consider the following:

- For MongoDB and MS SQL workloads, the alternate location must be discovered but should not be connected or configured.
- For Oracle workloads, the alternate location must be discovered and configured and not connected.

Before you begin

Ensure that the database is prepared for snapshots. For details review the plug-in configuration notes in the [Veritas CloudPoint documentation](#).

To configure applications for point-in-time recovery

- 1 Connect to the virtual machine that hosts the applications.
 - After the cloud assets are discovered, go the **Virtual Machines** tab.
 - Select the virtual machine where the application is hosted. On the top right, click **Connect VM**
 - Enter the credentials.
 - Click **Connect**.
 - After the virtual machines are connected, the virtual machines state is updated to **Configure**.

Note: For Microsoft SQL Server, you need perform this process manually. See, the Configuring the Windows-based on-host agent topic in the [Veritas CloudPoint documentation](#). After the next discovery cycle, the status of the virtual machine is updated to **Configure**.

- 2 Select the virtual machine where the application is hosted. On the top right, click **Configure application**.
- 3 After the process is complete, the application status is updated to configured.
- 4 The applications are displayed under the **Applications** tab after the next discovery.
- 5 Apply the protection plan. See the [NetBackup™ Web UI Backup Administrator's Guide](#).

To edit or update virtual machine credentials

- 1 Go to the **Virtual Machines** tab.
- 2 Select the virtual machines for which you want to update credentials. On the top right, click **Edit credentials**.
- 3 Update the credentials and click **Connect**.

To edit or update application configuration

- 1 Go to the **Applications** tab.
- 2 Select the application for which you want to update. On the top right, click **Edit configuration**
- 3 Update the credentials and click **Configure**.

Configure CloudPoint servers in NetBackup

You can now add a CloudPoint server using the NetBackup Web UI. Starting with 9.0, the CloudPoint can discover cloud assets on Amazon Web Services and Microsoft Azure US Government cloud.

Consider the following important points:

- You can associate multiple CloudPoint servers to a NetBackup master server. But you can associate only one CloudPoint server to one NetBackup master server.
- You can associate multiple media servers to a CloudPoint server. Only the media servers that are linked to your NetBackup master server can be linked to a CloudPoint server.
- You can now manage CloudPoint and control discovery of assets from the NetBackup WebUI, REST API, and CLI without interacting with the CloudPoint interfaces.

The following table describes the underlying tasks.

Table 3-5 Configuring CloudPoint servers

Task	Description
Add a CloudPoint server	To add a CloudPoint server in NetBackup, you must add the credentials and validate the certificate of the CloudPoint server. See “Add a CloudPoint server” on page 31.

Table 3-5 Configuring CloudPoint servers (*continued*)

Task	Description
Add cloud providers	To discover assets on the CloudPoint server, you must add the cloud providers. See “Add a cloud provider for a CloudPoint server” on page 32.
Discover assets on CloudPoint server	You can discover assets on the CloudPoint server. See “Discover assets on CloudPoint server” on page 35.
Associate media servers	To offload snapshots and restore workflows to a media server, you must associate the media server to the CloudPoint server. See “Associate media servers with a CloudPoint server” on page 35.

Configure a third-party CA certificate

You can use a self-signed or a third-party certificate to validate your CloudPoint server.

Consider the following points:

- For Windows, you can give a certificate as a file path or install the third party certificate in the Trusted Root Certificates authorities.
- To switch from a self-signed certificate to a third-party certificate for an already added CloudPoint server, you can update the `tpconfig` command or edit the CloudPoint server API or from NetBackup WebUI.

To configure a third-party CA certificate

- 1 Generate the third party certificate and private key for your CloudPoint server.
- 2 Run the `./cloudpoint/scripts/cp_certificate_management.sh` script to upload your certificate and keys to the CloudPoint server.
- 3 In NetBackup, create a certificate file and append the certificate of root and all intermediate CAs in the pem file.
- 4 In the `bp.conf` file, create the following entries:
 - `ECA_TRUST_STORE_PATH = /certificate.pem`
 - (Optional) `VIRTUALIZATION_CRL_CHECK = CHAIN`
 - (Optional) `ECA_CRL_PATH = /crls`

Note:

- The ECA_CRL_PATH option specifies the path to the directory where the Certificate Revocation Lists (CRL) of the external certificate authority (CA) are located. All files in ECA_CRL_PATH must be in pem format.
 - VIRTUALIZATION_CRL_CHECK option is only required if you want to check the revocation status of the certificate. By default, the VIRTUALIZATION_CRL_CHECK option is disabled.
 - You can disable, LEAF, or CHAIN the value of the VIRTUALIZATION_CRL_CHECK option. For LEAF, revocation status of the leaf certificate is validated against the CRL. For CHAIN, revocation status of all certificates from the certificate chain are validated against the CRL.
- 5** Add the CloudPoint server to NetBackup or run the `tpconfig` command to update the certificate for a CloudPoint server already added to NetBackup.

Note: Following should be the order in which the certificates are uploaded:

- Leaf
- Intermediate
- Root

If the certificates are not uploaded in the correct order, the CloudPoint might not work.

Add a CloudPoint server

You can add a CloudPoint server using NetBackup. You must provide the CloudPoint server credentials and validate the certificate.

To add a CloudPoint server

- 1** On the left, click **Cloud**.
- 2** Click on the **CloudPoint server** tab.
- 3** Click **Add**.
- 4** In the **CloudPoint server** field, enter one of the following:
 - The host name or IP address of the CloudPoint server.

The host name or IP address must be the same as the one you have provided at the time of CloudPoint configuration during CloudPoint installation.

- If the DNS server is configured, enter the FDQN of the CloudPoint server.
- 5** In **Port** field, enter the port number for the CloudPoint server.
The default port value is 443.
 - 6** Click **Validate**.
 - 7** In the **Validate certificate** dialog box, click **Accept**.
 - 8** Enter the CloudPoint server credentials.
 - 9** Click **Save**.

Add a cloud provider for a CloudPoint server

You can protect the assets on the Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure cloud providers. Starting with 9.0, the CloudPoint server can discover Amazon Web Services and Microsoft Azure US Government cloud workloads.

To add a cloud provider for CloudPoint server

- 1** On the left, click **Cloud**.
- 2** Click on the **CloudPoint server** tab.
- 3** Click the **Providers** tab or click **Add** under the cloud provider for which you want to add a configuration.
- 4** Enter a value in the **Configuration Name** field, in the **Add configuration** pane.
- 5** Select the preferred **CloudPoint server**.
- 6** Click **Save**.

7 Enter the required details.

Cloud provider	Parameter	Description
Microsoft Azure	Tenant ID	The ID of the AAD directory in which you created the application.
	Client ID	The application ID.
	Secret Key	The secret key of the application.
	Regions	One or more regions in which to discover cloud assets. Note: If you configure a government cloud, select US Gov Arizona, US Gov Texas US, or Gov Virginia.
	Resource Group prefix	The string with which you want to append all the resources in a resource group.
	Protect assets even if prefixed Resource Groups are not found	The check box determines whether the assets are protected if they are not associated to any resource groups.
Amazon AWS Note: If the CloudPoint server is configured with IAM Config, the Access Key and Secret Key options are not available.	Access Key	The access key ID, when specified with the secret access key, authorizes CloudPoint to interact with the AWS APIs.
	Secret Key	The secret key of the application.
	Regions	One or more AWS regions in which to discover cloud assets. Note: If you configure a government cloud, select us-gov-east-1 or us-gov-west-1.

Cloud provider	Parameter	Description
Google Cloud Platform	Project ID	The ID of the project from which the resources are managed. Listed as in the <code>project_id</code> JSON file.
	Client Email	The email address of the Client ID. Listed as <code>client_email</code> in the JSON file.
	Private Key	The private key. Listed as <code>private_key</code> in the JSON file. Note: You must enter this key without quotes. Do not enter any spaces or return characters at the beginning or end of the key.
	Zones	A list of zones in which the provider operates.

- 8 Enter the connection and authentication details in the **Add Configuration** pane.
- 9 Click **Save**.

The assets on the cloud providers are automatically discovered.

IAM Role for AWS Configuration

If the snapshot management server (CloudPoint) is deployed in cloud, AWS configuration can be configured to use IAM role for authentication.

See [“Add a cloud provider for a CloudPoint server”](#) on page 32.

Before proceeding, ensure the following:

- IAM role is configured within AWS. See the *NetBackup CloudPoint Install and Upgrade Guide* for details.
- After you upgrade NetBackup and CloudPoint to the latest version, you need to update the credentials. Run the `tpconfig -update` command.

Note: Post upgrade, credentials are updated to support only IAM role.

The following implementations of IAM role are supported:

- **Source account:** In this case, the cloud assets that need to be protected are in the same AWS account as CloudPoint. Thus, AWS cloud is aware of the AWS account ID and role name, you need to only select the region.
- **Cross account:** In this case, the cloud assets that need to be protected are in a different AWS account than CloudPoint. Thus, you need to enter the target account and the target role name details along with the region so that CloudPoint can access those assets.

You need to establish a trust relationship between the source and the target account. For more details, refer to the *Across AWS Accounts Using IAM Roles* related information in the *Amazon Web Services* documentation.

Associate media servers with a CloudPoint server

You can use a media server to offload the snapshots and restores jobs of your cloud. To enable that you must associate one or more media servers to a CloudPoint server. The media servers must be in an active state to run the snapshot or restore jobs. The media server that you associate with the CloudPoint server must be associate to your NetBackup master server also. However, the discovery jobs run on the NetBackup master server only.

To associate media servers with a CloudPoint server

- 1 On the left, click **Cloud**.
- 2 Click on the **CloudPoint server** tab.
- 3 From the menu next to the CloudPoint server, click **Advanced settings**.
- 4 In the **Advance settings** dialog box, select one or more media servers that you want associate with the CloudPoint server.
- 5 Click **Save**.

Discover assets on CloudPoint server

After you configure your cloud providers to a CloudPoint server, you can discover the assets and assign protection plans. As part of this operation, first the cloud discovery is triggered on CloudPoint server. The CloudPoint server discovers all the assets from the Cloud. After the discovery on CloudPoint server completes, the assets in NetBackup are updated with the CloudPoint server assets. If you disable a CloudPoint server, all the assets associated with that server are no longer protected.

Note: The timeout for CloudPoint discovery is 30 minutes. If the discovery on CloudPoint server takes more than 30 minutes, the first discovery operation times out. But the second operation is continues which updates the NetBackup assets with the CloudPoint server assets.

To discover assets on CloudPoint server

- 1 On the left, click **Cloud**.
- 2 Click on the **CloudPoint server** tab
- 3 From the menu next to the CloudPoint server, click **Discover**.

Edit a CloudPoint server

You can update the CloudPoint server credentials. However, you cannot edit the Host name, IP address or Port of a CloudPoint server.

To edit a CloudPoint server

- 1 On the left, click, **Cloud**.
- 2 Click on the **CloudPoint server** tab.
- 3 From the menu next to the CloudPoint server, click **Edit**.

You can only edit the credentials for CloudPoint Server. You must validate the certificate before you can update the credentials.

- 4 Update the credentials.
- 5 Click **Save**.

Enable or disable a CloudPoint server

Based on your preference, you can enable or disable a CloudPoint server. If you disable a CloudPoint server, you cannot discover assets or assign protection plans.

To enable or disable a CloudPoint server

- 1 On the left, click **Cloud**.
- 2 Click on the **CloudPoint server** tab.
- 3 Based on the CloudPoint server status, select **Enable** or **Disable**.

Recovering cloud assets

This chapter includes the following topics:

- [Recover a cloud asset to its original location](#)
- [Recover a cloud asset to an alternate location](#)
- [Perform rollback recovery of cloud assets](#)

Recover a cloud asset to its original location

For Google Cloud Platform, when you restore a snapshot that was created before the upgrade, if the source disk is not present, a default restored disk, pd-standard is created.

To recover a cloud asset to its original location

- 1 On the left, click **Cloud**.
- 2 Depending on the cloud asset type, click on the **Virtual Machines**, **Applications**, or **Volumes** tab.

All the discovered cloud assets for the respective category are displayed.

- 3 Double-click on the protected asset that you want recover.
- 4 Click the **Recovery points** tab.

The available images are listed in rows with the backup timestamp for each image.

- 5 On the top right for the preferred recovery point, select **Original location**.
- 6 Click **Start recovery**.
- 7 On the left, click **Activity monitor** to view the job status.

Recover a cloud asset to an alternate location

Considerations

- Cloud assets for Google Cloud Platform cannot be restored to an alternate location.
- To restore a replicated copy of EC2 instance to an alternate location, the key-pair names must be same on the source and destination region. If not, create a new key-pair in the destination region that is consistent with the key-pair in the source region.

To recover a cloud asset to an alternate location

- 1 On the left, click **Cloud**.
- 2 Depending on the cloud asset type, click on the **Virtual Machines**, **Applications**, or **Volumes** tab.
All the discovered cloud assets for the respective category are displayed.
- 3 Double-click on the protected asset you want recover.
- 4 Click the **Recovery points** tab.
The available images are listed in rows with the backup timestamp for each image.
- 5 On the top right for the preferred recovery point, select **Alternate location**.
- 6 Select the location where you want to restore the cloud asset.
- 7 Click **Start Recovery**.
- 8 In the left, click **Activity monitor** to view the job status.

Perform rollback recovery of cloud assets

The rollback recovery of a cloud asset overwrites the existing data on the original asset. Unlike original or alternate location restore, a new copy is not created of the restored image, but the existing data on the source is replaced.

Note: Snapshot replicas cannot be rolled back.

To perform rollback recovery of the cloud asset

- 1 On the left, click **Cloud**.
- 2 For supported cloud asset type, click on the **Virtual Machines**.
All the discovered cloud assets for the respective category are displayed.

- 3** Double-click on the protected asset you want to recover.
- 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, click **Recover > Rollback**.
- 6** Click **Start recovery**. The existing data is overwritten.
- 7** On the left, click **Jobs** to view the job status.

Troubleshooting protection and recovery of cloud assets

This chapter includes the following topics:

- [Troubleshoot cloud workload protection issues](#)

Troubleshoot cloud workload protection issues

Review the following log files to troubleshoot any issues with protection of cloud assets:

- [Log files for configuration](#)
- [Log files for snapshot creation](#)
- [Log files for restore operations](#)
- [Log files for snapshot deletion](#)

During troubleshooting, ensure that you have also reviewed the limitations. See [“Limitations and considerations”](#) on page 18.

For troubleshooting issues, see the [NetBackup Status Codes Reference Guide](#).

To view the CloudPoint log files, see the CloudPoint logs topic in the *Veritas NetBackup CloudPoint Install and Upgrade Guide*.

Log files for configuration

Use the following logs to troubleshoot cloud configuration issues.

Table 5-1 Log files for configuration

Process	Logs
<p>tpconfig</p> <p>tpconfig command is one way for registering CloudPoint in NetBackup.</p>	<p>Windows</p> <p><i>NetBackup install path/volmgr/debug/tpcommand</i></p> <p>UNIX</p> <p><i>/usr/opensv/volmgr/debug/tpcommand</i></p>
<p>nbwebservice</p> <p>Plug-ins are configured using NetBackup REST API.</p>	<p>Windows</p> <p><i>NetBackup install path/webserver/logs</i></p> <p>UNIX</p> <p><i>/usr/opensv/wmc/webserver/logs</i></p> <p><i>/usr/opensv/logs/nbwebservices</i></p>
<p>nbemm</p> <p>nbemm stores the CloudPoint server and plug-in information in EMM database</p>	<p>Windows</p> <p><i>NetBackup install path/path/logs/nbemm</i></p> <p>UNIX</p> <p><i>/usr/opensv/logs/nbemm</i></p>

Log files for asset discovery

Use the following logs to troubleshoot asset discovery issues.

Table 5-2 Log files for asset discovery

Process	Logs
<p>ncfnbcs</p> <p>Verifies if discovery was completed or not.</p>	<p>Windows</p> <p><i>NetBackup install path/bin/vxlogview -o 400</i></p> <p>UNIX</p> <p><i>/usr/opensv/netbackup/bin/vxlogview -o 400</i></p>
<p>Picloud</p> <p>Provides the details of discovery operation.</p>	<p>Windows</p> <p><i>NetBackup install path/bin/vxlogview -i 497</i></p> <p>UNIX</p> <p><i>/usr/opensv/netbackup/bin/vxlogview -i 497</i></p>

Table 5-2 Log files for asset discovery (*continued*)

Process	Logs
<p>nbwebservice</p> <p>To get details about the asset database workflows that are part of the discovery operation.</p> <p>Note: Refer to the same log files for details of assets that are added to protection plan.</p>	<p>Windows</p> <p><i>NetBackup install path/webserver/logs</i></p> <p>UNIX</p> <p><i>/usr/opensv/wmc/webserver/logs</i></p> <p><i>/usr/opensv/logs/nbwebservicess</i></p>

Log files for snapshot creation

Use the following logs to troubleshoot snapshot creation issues.

Table 5-3 Log files for snapshot creation

Process	Logs
<p>nbpem</p> <p>nbpem PID for given job is available in the NetBackup activity monitor.</p>	<p>Windows</p> <p><i>NetBackup install path/bin/vxlogview -o 116</i></p> <p>UNIX</p> <p><i>/usr/opensv/netbackup/bin/vxlogview -o 116</i></p>
<p>nbjm</p> <p>nbjm PID for given job is available in the NetBackup activity monitor.</p>	<p>Windows</p> <p><i>NetBackup install path/bin/vxlogview -o 117</i></p> <p>UNIX</p> <p><i>/usr/opensv/netbackup/bin/vxlogview -o 117</i></p>
<p>nbcs</p> <p>nbcs PID for given job is available in the NetBackup activity monitor.</p>	<p>Windows</p> <p><i>NetBackup install path/bin/vxlogview -i 366 -P nbcs_process_id</i></p> <p>UNIX</p> <p><i>/usr/opensv/netbackup/bin/vxlogview -i 366 -P nbcs_process_id</i></p> <p>The nbcs logs are available at the following location:</p> <p>Windows</p> <p><i>NetBackup install path/logs/ncfnbcs</i></p> <p>UNIX</p> <p><i>/usr/opensv/logs/ncfnbcs</i></p>

Table 5-3 Log files for snapshot creation (*continued*)

Process	Logs
nbrb nbrb is requested to provide a media server for a given job. For Cloud, a particular media server is picked up from the associated list of media servers for a CloudPoint server.	Windows <i>NetBackup install path/bin/vxlogview -o 118</i> UNIX <i>/usr/opensv/netbackup/bin/vxlogview -i 118</i>

Log files for restore operations

Use the following logs to troubleshoot restore issues.

Table 5-4

Process	Logs
nbwebservice The snapshot restore operation is triggered by NetBackup REST API.	Windows <i>NetBackup install path/webserver/logs</i> UNIX <i>/usr/opensv/wmc/webserver/logs</i> <i>/usr/opensv/logs/nbwebservices</i>
bprd The NetBackup REST API communicates with bprd to initiate restore	Windows <i>NetBackup install path/netbackup/logs</i> UNIX <i>/usr/opensv/netbackup/logs/bprd</i>
ncfnbcs nbcs PID for given job is available in the NetBackup activity monitor.	Windows <i>NetBackup install path/bin/vxlogview -i 366 -P nbcs_process_id</i> UNIX <i>/usr/opensv/netbackup/bin/vxlogview -i 366 -P nbcs_process_id</i>

Log files for snapshot deletion

Use the following logs to troubleshoot snapshot deletion issues.

Table 5-5 Log files for snapshot deletion

Process	Logs
<p>bpdm</p> <p>The snapshot delete or clean-up operation is triggered by bpdm.</p>	<p>Windows</p> <p><i>NetBackup install path/netbackup/logs</i></p> <p>UNIX</p> <p><i>/usr/opensv/netbackup/logs/bpdm</i></p>
<p>ncfnbcs</p> <p>nbcs PID for given job is available in the NetBackup activity monitor.</p>	<p>Windows</p> <p><i>NetBackup install path/bin/vxlogview -i 366 -P nbcs_process_id</i></p> <p>UNIX</p> <p><i>/usr/opensv/netbackup/bin/vxlogview -i 366 -P nbcs_process_id</i></p>

Performing granular restore

This chapter includes the following topics:

- [About granular restore](#)
- [Supported environment list](#)
- [List of supported file systems](#)
- [Before you begin](#)
- [Limitations and considerations](#)
- [Restoring files and folders cloud virtual machines](#)
- [Restoring volumes on cloud virtual machines](#)
- [Troubleshooting snapshot restore process for Microsoft Azure-specific cloud](#)

About granular restore

NetBackup enables you to perform a granular restore of files and folders on cloud virtual machines. You can create snapshots and restore, at the same time you can also locate and restore individual files and folders. You can also restore volumes from virtual machines.

This process is known as granular restore in which each single file in the snapshot is considered as a granule or more commonly referred to as single file restore. NetBackup makes an inventory of all the files within a snapshot using an indexing process. You can restore specific files from a snapshot only if that snapshot has been indexed by NetBackup.

The following table helps you understand the flow of enabling granular restore of volumes, files, and folders:

Table 6-1 Granular restore tasks

Task	Description
Connect virtual machines	Connect the virtual machines that you want to use to perform granular restore.
Discover assets on virtual machine	Use the Discover option. Navigate to Cloud > CloudPoint servers > CloudPoint server > Actions > Discover .
Create protection plan	Create a protection plan. Ensure that the Enable granular recovery for files or folders check box is selected in the Backup options of the protection plan.
Subscribe discovered assets to the protection plan	Add the assets on the VMs connected in the previous step to the protection plan that has the indexable attribute enabled granular restore.
Execute protection plan	Schedule backup job and indexing or use the Backup now option. The backup job starts immediately.
Restore file or folder or Restore volumes	Perform granular restore of a file, folder, or volume.

Supported environment list

The following table lists the supported versions.

Table 6-2 Supported versions

Application	Version
NetBackup	9.0
NetBackup backup host OS	RHEL 7.x
CloudPoint host OS	<ul style="list-style-type: none"> ■ RHEL 7x and later ■ Ubuntu 18.04 LTS and 16.04 LTS

Table 6-2 Supported versions (*continued*)

Application	Version
Cloud providers	<ul style="list-style-type: none"> ■ Amazon Web Services ■ Microsoft Azure ■ Google Cloud Platform <p>Note: Granular restore is not supported on Windows environment on Google Cloud Platform.</p>
CloudPoint or agent instance type	<ul style="list-style-type: none"> ■ Amazon AWS: t2.large/t3.large ■ Microsoft Azure: D2s_V3Standard ■ Google Cloud Platform: n1.Standard2 and larger
CloudPoint agent host to be protected	<ul style="list-style-type: none"> ■ Linux OS: RHEL 7.7 and 7.6 ■ Windows OS Version: 2016 and 2012 <p>Note: Granular restore is not supported for RHEL 8.X platform.</p>

List of supported file systems

The following table provides details about supported files systems.

Platform	Discovered file system	Partition layouts
RHEL (With consistent snapshot property)	<ul style="list-style-type: none"> ■ ext3 ■ ext4 <p>Note: Granular restore is not supported for RHEL 8.X platform.</p>	<ul style="list-style-type: none"> ■ GPT ■ MBR ■ No layout (direct FS)
Windows (With consistent snapshot property)	NTFS	<ul style="list-style-type: none"> ■ GPT ■ MBR

Note: Consistent snapshot is not supported for ext2 file system version.

Before you begin

Ensure the following points are addressed before you perform granular restore.

- Configured CloudPoint server and VM to be protected with granular restore enabled have the following requirements:
- Microsoft Azure: The CloudPoint host and the connected VM must be in the same subscription and region.
- Amazon AWS: The CloudPoint host and the connected VM must be in the same account and region.
- Google Cloud Platform: The CloudPoint host and the connected VM must be in the same project.
- The host must be in a connected state and must have required supported configuration.
- The cloud plug-in must be configured to protect the assets in the region in which the CloudPoint host is deployed.
- The host must be in a connected state and must have required supported configuration.
- The host must have the **fsConsistent** and **indexable** flags enabled when connected. The **fsConsistent** flag allows the files systems on the host to be snapped by the CloudPoint and indexable flag allows the host to be indexed. The **indexable** flag can be set as true only if the **fsConsistent** flag is set as true.
- Protection plan must have the **Enable Granular restore for files and folders** check box enabled.
- Apart from the boot disk and disk that is mounted on "/cloudpoint", no extra disk should be attached to CloudPoint instance explicitly.
- File systems on the host must be supported. See "[List of supported file systems](#)" on page 47.
- Configure port 5671 and 443 for open CloudPoint host.
- For an agentless restore, port 22 must be configured on the indexable virtual machines for agent connection.
- For on-host restore, port 5671 and 3389 (RDP) must be open on the target virtual machine for agent connection. RDP is only for used for configuration and is optional.
- Ensure that you have appropriate permissions to perform a granular restore. See the Role permissions topic in the *NetBackup Web UI Administrator's Guide*.
- If you want to restore volume to same the virtual machine and location, you must detach existing volume and free the slot and then try to restore.

Limitations and considerations

Consider the following important points for granular restore.

- After a restore job is completed, you cannot expand the directories in the **File List** section of the restore job.
- If adequate space is not available on the target location, the restore operation fails before the copy operation begins.
- In the activity monitor summary, when the restore job starts it shows the current file which is the first entry in the restore items. After the job is complete, the summary goes blank.
- Bytes transferred and estimated bytes in activity monitor are not updated and shown as 0.
- The maximum number of indexing jobs that CloudPoint supports is limited based on the following conditions:
 - The number of attachment points for a data disk available on the CloudPoint host minus 1 and the instance type. CloudPoint metadata volume consumes this one attachment point.
 - The resource availability of the CPU or memory of the CloudPoint machine.
- The ephemeral storage devices like Amazon AWS instance store volumes and Microsoft Azure temporary disks are ignored when a snapshot is performed. These devices are also ignored for indexing as well.
- The file systems that are created on LVM or LDM disks are ignored when host consistent snapshots are created and indexed.
- If any unsupported file systems are present on the host, the host can't be added to the protection plan that is created for granular restore. The protection plans for granular restore have the **Enable granular recovery for files or folders** check box value set to true.
- CloudPoint communicates the number of index jobs that can be run to NetBackup. NetBackup then throttles the requests. By default, the number of index jobs is initialized to 2. Post discovery of CloudPoint host capabilities, it is increased to number of disk slots available. However, you can update the value for indexing `max_jobs=<value>` in `flexsnap.conf` file to override this limit.
- CloudPoint host limits the number of disk slots that the cloud providers enforce. NetBackup throttles the indexing requests to CloudPoint. To achieve this request, during Cloud Asset discovery process, NetBackup fetches CloudPoint host capabilities. These capabilities include the **Max no of index jobs** parameter. This parameter is used to limit the requests that are sent to CloudPoint and index job queue in NetBackup. By default, the maximum number of parallel

indexing jobs is 2. But once the cloud plug-in is configured which discovers the CloudPoint host, the capability API fetches the number of max jobs based on attachment points and available resources. You can set the limit by adding the `indexing max_jobs=x` entry in the **config** file of the CloudPoint host. If the CloudPoint host receives number of indexing requests more than its capability, the requests are queued.

- When an indexing operation is in progress, if any OS errors occur while crawling files, directories, or other entries, the errors are ignored and the indexing operation continues. To restore the missing files, you must initiate the granular restore operations on the parent folder.

- If a mount point is not visible in the tree on the left panel for browsing when you add files or folders from the recovery point, it can be because of the following reasons:

- The "/" (root file system) is on an LVM, and

- The mount point is not directly related to "/" (root file system)

In this scenario, search for the mount point from the right panel and then restore the files or folders successfully.

For example, if a disk is mounted on /mnt1/mnt2 where /mnt1 is any directory on the "/" (root FS which is on LVM setup) and mnt2 is a mount point inside mnt1, the "mnt2" is not visible in the tree on the left panel. However, you can search and restore files or folders inside mount point.

- Application and disk restore is supported from RHEL 8.2 to RHEL 8.2 target only.
- While restoring application or file systems from one OS version to another OS version, refer to the OS and application vendor's compatibility matrix.
- Granular restore is not supported for RHEL 8.x platform.
- While restoring a drive as source to an alternate folder as destination, the user group cannot perform the write operation on the newly created folder due to lack of the write permission.

Restoring files and folders cloud virtual machines

You can restore a single file or folder from a cloud virtual machine.

Note: For Microsoft Azure and Google Cloud Platform, NetBackup supports snapshot and recovery of cloud assets that are encrypted using the keys are the manager provides.

To restore a file or folder

- 1** On the left, click **Cloud**.
- 2** Click on the **Virtual machines** tab.
- 3** Select the virtual machine where the application is hosted. On the top right, click **Connect VM**.
- 4** After the VM is connected, on the top right, click **Add protection**.
- 5** Select a protection plan that is created for granular recovery of files and folders and click **Next**.
- 6** Click **Protect**.
- 7** To execute the protection plan, click **Backup now**.
- 8** After a snapshot and the two indexing job for the assets are complete, to view the recovery points, click the **Recovery points** tab.
- 9** On the top right for the preferred recovery point, select **Restore files and folders**.

You can also apply date filters to search for across recovery points. In case of replication, you click **Recover** and then select Restore files and folders.

- 10** In the Add file step, click **Add**.
- 11** In the **Add files and folders** dialog box, select the files you want to restore and click **Add**.

You can click the folders or drives on the left to expand and view the files in a particular folder. You can search files based on their names or extensions.

- 12** Click **Next**.
- 13** From the **Target VM** list in the Recovery target step, select a VM.

A list with all connected VMs having same operating systems as original target host is displayed. If you do not select a VM, the files are restored to the original VM.

- 14** In the **Files restore** options, select one of the following options:
 - **Restore everything to original directory**
 - **Restore everything to a different directory**
 You must then provide a directory location. You can also enter a UNC path to the location.

- 15** Click **Next**.
- 16** In the Recovery options step, select the preferred option:
 - **Append string to file names**

In the **String** field, enter the string that you want use to append. The string is appended before the last extension of a file.

- **Overwrite existing files**
You must have appropriate permissions.
- (If you selected the **Restore everything to a different directory** option)
Create new files for hard links

17 Click **Next**.

18 In the Review step, view the selected options and click **Start Recovery**.

The restore job for the selected files is triggered. You can view the job details on the Activity monitor. After the job is successful, you can see summary of restored files in the job details.

Restoring volumes on cloud virtual machines

You can restore one or more volumes on a virtual machine.

To restore a volume

- 1** On the left, click **Cloud**.
- 2** Click on the **Virtual machines** tab.
- 3** Select the virtual machine where the application is hosted.
- 4** After the VM is connected, on the top right, click **Add protection**.
- 5** Select a protection plan and click **Next**.
- 6** Click **Protect**.
- 7** To execute the protection plan, click **Backup now**.
- 8** To view the recovery points, click the **Recovery points** tab.
- 9** On the top right for the preferred recovery point, select **Restore volumes**.
You can also apply date filters to search across the recovery points.
- 10** In the **Restore volumes** dialog box, select one or more volumes.

- 11 From the **Target VM** list, select the VM on which you want to restore the volumes.

In case of restore from a replicated (non-primary) VM, the restore to original location is not supported. If you do not select a VM, the files are restored to the original VM.

- 12 Click **Restore**.

The restore job for the selected volumes is triggered. You can view the job details on the Activity monitor.

Troubleshooting snapshot restore process for Microsoft Azure-specific cloud

When you trigger a subsequent (twice) restore operation on the same VM, an error occurs during restore operation. This error causes the following issues:

- The tags from original OS disk are not copied to newly created restored OS disk.
- User login might fail after the VM restore due to ssh failure.

Workaround:

Check if the ssh daemon is running on the system. If not, then perform the steps that are mentioned in the

<https://docs.microsoft.com/en-us/azure/virtual-machines/troubleshooting/troubleshoot-ssh-connection> topic.