

NetBackup™ Web UI Cloud Administrator's Guide

Release 10.0

Last updated: 2022-03-27

Legal Notice

Copyright © 2022 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	7
	7
	About the NetBackup web UI	7
	Terminology	8
	Sign in to the NetBackup web UI	10
	Sign out of the NetBackup web UI	12
Chapter 2	Monitoring NetBackup	13
	The NetBackup dashboard	13
	Job monitoring	13
	Search for or filter jobs in the jobs list	14
Chapter 3	Managing and protecting cloud assets	16
	About protecting cloud assets	17
	Limitations and considerations	18
	Configure CloudPoint servers in NetBackup	19
	Configure a third-party CA certificate	20
	Add a CloudPoint server	22
	Add a cloud provider for a CloudPoint server	22
	Associate media servers with a CloudPoint server	26
	Discover assets on CloudPoint server	27
	Edit a CloudPoint server	28
	Enable or disable a CloudPoint server	29
	(Optional) Add the CloudPoint extension	29
	Managing intelligent cloud groups	29
	Create an intelligent cloud group	30
	Delete an intelligent cloud group	33
	Protecting cloud assets or intelligent cloud groups	34
	Customize or edit protection for cloud assets or intelligent groups	
	36
	Remove protection from cloud assets or intelligent groups	37
	Cloud asset cleanup	37
	AWS and Azure government cloud support	38

About protecting Microsoft Azure resources using resource groups	39
Before you begin	39
Limitations and considerations	40
About resource group configurations and outcome	40
Troubleshoot resource group permissions	43
About the NetBackup Accelerator for cloud workloads	44
How the NetBackup Accelerator works with virtual machines	44
Accelerator forced rescan for virtual machines (schedule attribute)	45
Accelerator backups and the NetBackup catalog	46
Accelerator messages in the backup job details log	46
Configuring backup schedule for cloud workloads	47
Backup options for cloud workloads	49
Snapshot replication	52
Configure AWS snapshot replication	52
Using AWS snapshot replication	55
Support matrix for account replication	57
Protect applications in-cloud with application consistent snapshots	59
Discovering PaaS assets	61
Chapter 4	
Recovering cloud assets	62
Recovering cloud assets	62
Perform rollback recovery of cloud assets	68
Recovering PaaS assets	69
Chapter 5	
Performing granular restore	72
About granular restore	72
Supported environment list	73
List of supported file systems	74
Before you begin	75
Limitations and considerations	76
Restoring files and folders from cloud virtual machines	78
Restoring volumes on cloud virtual machines	79
Performing steps after volume restore containing LVM	80
Troubleshooting	82

Chapter 6	Troubleshooting protection and recovery of cloud assets	84
	Troubleshoot cloud workload protection issues	84
	Troubleshoot PaaS workload recovery issues	88

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 primary 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 for workload protection.
- Protection of assets is achieved through protection plans, job management, and visibility of the protection status of assets.

- Workload administrators can create protection plans, manage credentials, subscribe assets to the protection plans that meet the SLO, monitor protection status, and perform self-service recovery of assets.

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 features that the user can access in the web UI. For example, access to any workload assets, protection plans, or credentials.
- 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, except for the Enhanced Auditing (EA).

Monitor NetBackup jobs

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

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

Protection plans offer the following benefits:

- A default workload administrator can select the protection plans to use to protect assets.
- With the necessary RBAC permissions, a workload administrator can create and manage protection plans, including the backup schedules and storage that is used.
- When you select from your available storage, you can see any additional features available for that storage.

Self-service recovery

The NetBackup web UI makes it easy for a workload administrator to recover VMs, databases, or other asset types applicable to that workload.

Terminology

The following table describes the concepts and terms in web user interface.

Table 1-1 Web user interface terminology and concepts

Term	Definition
Asset group	See <i>intelligent group</i> .
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.
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>These groups appear under the tab Intelligent VM groups or Intelligent groups.</p>
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. The role administrator 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.</p>
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.

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

Term	Definition
Workload	The type of asset. For example: VMware, RHV, AHV, Microsoft SQL, Oracle, Cloud, or Kubernetes.

Sign in to the NetBackup web UI

Authorized users can sign in to a NetBackup primary server from a web browser, using the NetBackup web UI.

For more information, refer to the *Authorized users* section in the *NetBackup™ Web UI Administrator's Guide*.

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 primary server using a username and password

- 1 Open a web browser and go to the following URL.
`https://primaryserver/webui/login`
The *primaryserver* is the host name or IP address of the NetBackup primary server that you want to sign in to.
- 2 Depending on the sign-in options that are available, choose from the following:
 - Enter your credentials and click **Sign in**.
 - If the default method is not username and password, click **Sign in with username and password**. Then enter your credentials.

The following are example credentials:

For this type of user	Use this format	Example
Local user	<i>username</i>	jane_doe
Windows user	<i>DOMAINusername</i>	WINDOWS\jane_doe

For this type of user	Use this format	Example
UNIX user	<i>username</i>	john_doe

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://primaryserver/webui/login>
 The *primaryserver* is the host name or IP address of the NetBackup primary 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 primary server using SSO

- 1 Open a web browser and go to the following URL.
<https://primaryserver/webui/login>
 The *primaryserver* is the host name or IP address of the NetBackup primary 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 primary 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 NetBackup

This chapter includes the following topics:

- [The NetBackup dashboard](#)
- [Job monitoring](#)
- [Search for or filter jobs in the jobs 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.

Job monitoring

Use the **Jobs** node in the Activity monitor to monitor the jobs in your NetBackup environment. The default view for jobs is the **List view** that contains the non-hierarchical list of all the jobs. You can also use the **Hierarchical view** to see the hierarchy of parent and child jobs.

List view

Hierarchy view

<input type="checkbox"/>	Job ID ↑	Type	Client or display name	Job state
<input type="checkbox"/>	22322314	Backup	pe...10	Done
<input type="checkbox"/>	22322315	Backup	pe...10	Done
<input type="checkbox"/>	22322316	Backup	pe...10	Done
<input type="checkbox"/>	22322317	Backup	pe...10	Done
<input type="checkbox"/>	22322318	Backup	pe...10	Done
<input type="checkbox"/>	22322319	Backup	pe...08	Done

<input type="checkbox"/>	Job ID ↑	Type	Client or display name	Job state
▼ <input type="checkbox"/>	22322314	Backup	pe...10	Done
<input type="checkbox"/>	22322315	Backup	pe...10	Done
<input type="checkbox"/>	22322316	Backup	pe...10	Done
<input type="checkbox"/>	22322317	Backup	pe...10	Done
<input type="checkbox"/>	22322318	Backup	pe...10	Done
▼ <input type="checkbox"/>	22322319	Backup	pe...08	Done
<input type="checkbox"/>	22322320	Backup	pe...08	Done
<input type="checkbox"/>	22322321	Backup	pe...08	Done
<input type="checkbox"/>	22322322	Backup	pe...08	Done
<input type="checkbox"/>	22322323	Backup	pe...08	Done

The type of jobs that you can view and manage depend on the RBAC role that you have. For example, a workload administrator (such as the Default VMware Administrator role) can view and manage only jobs for that workload. In contrast, the Administrator role lets you view and manage all NetBackup jobs.

If you have an RBAC role that allows access to jobs, you can see the jobs list in the job hierarchy view. For example, the Default VMware Administrator role lets you see VMware jobs in the hierarchy view. However, if you only have access to one or more VMs (asset-level access), no jobs display in the job hierarchy view.

Search for or filter jobs in the jobs list

You can search for jobs in the Activity monitor or create filters to customize the jobs that are displayed.

Search for jobs in the jobs list

The search feature lets you search for the following job information: status code (complete status code #); policy name; client or display name; client; job ID (complete job ID #), or the job's parent ID.

Search for jobs in the jobs list

- 1 Click **Activity monitor > Jobs**.
- 2 In the **Search** box, type the keyword you want to find. For example, a client name or a status code number.

Filter the job list

To filter the job list

- 1 Click **Activity monitor > Jobs**.
- 2 In the toolbar, click the **Filter** icon.
- 3 Click on a filter that you created. Or, click **All jobs** to display all of the available jobs.

Managing and protecting cloud assets

This chapter includes the following topics:

- [About protecting cloud assets](#)
- [Limitations and considerations](#)
- [Configure CloudPoint servers in NetBackup](#)
- [Managing intelligent cloud groups](#)
- [Protecting cloud assets or intelligent cloud groups](#)
- [Cloud asset cleanup](#)
- [AWS and Azure government cloud support](#)
- [About protecting Microsoft Azure resources using resource groups](#)
- [About the NetBackup Accelerator for cloud workloads](#)
- [Configuring backup schedule for cloud workloads](#)
- [Backup options for cloud workloads](#)
- [Snapshot replication](#)
- [Configure AWS snapshot replication](#)
- [Using AWS snapshot replication](#)
- [Support matrix for account replication](#)
- [Protect applications in-cloud with application consistent snapshots](#)

- [Discovering PaaS assets](#)

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 8.3, CloudPoint can now protect assets in AWS, Azure, Azure Stack hub and GCP clouds.

The following table describes the tasks.

Table 3-1 Configuring protection for cloud assets

Task	Description
Before you begin ensure that you have the appropriate permissions.	<p>To manage and protect cloud assets in the web UI you must have the workload administrator role or similar permissions. The NetBackup security administrator can manage your role permissions at an individual asset level or at the account or subscription level, or at a cloud provider level.</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>
Deploy CloudPoint	<p>Install CloudPoint in your environment.</p> <p>See "Add a CloudPoint server" on page 22.</p> <p>Review CloudPoint and NetBackup limitations.</p> <p>See "Limitations and considerations" on page 18.</p>
Configure the CloudPoint server using the NetBackup Administration Console	<p>Register the CloudPoint server in NetBackup.</p> <p>See, <i>NetBackup Snapshot Client Administrator's Guide</i>.</p>
Add a configuration	<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 22.</p> <p>For Amazon, you can choose to use IAM role.</p> <p>See "IAM Role for AWS Configuration" on page 26.</p>

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

Task	Description
Asset discovery	<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 “Discover assets on CloudPoint server” on page 27.</p>
Create a protection plan	<p>Create a 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 AWS snapshot replication” on page 52.</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 59.</p>
Recover cloud assets	<ul style="list-style-type: none"> ■ You can recover the assets using the recovery points. See “Recovering cloud assets” on page 62. See “Perform rollback recovery of cloud assets ” on page 68. ■ You can also restore the assets using the <code>nbcloudrestore</code> CLI utility. <p>Note: Do not use the <code>bprestore</code> CLI for restores</p> <p>See the NetBackup Commands Reference Guide.</p>
Troubleshooting	<p>See “Troubleshoot cloud workload protection issues” on page 84.</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.
- 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.
- By default, CloudPoint is configured with port 443.
- 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.
- For CloudPoint server, enhanced auditing is not supported. Thus, when you add or update a CloudPoint server, with non-root but NetBackup Admin rights, during auditing the user is shown as root.
- If you deploy CloudPoint using the CloudFormation template, when you register the on-host agent with the CloudPoint node using the command, the IP address used must be private IP and not public IP.

Configure CloudPoint servers in NetBackup

You can add a CloudPoint server using the NetBackup Web UI. Starting with 8.3, the CloudPoint server 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 primary 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 primary 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.

- For backup from snapshot jobs, the NetBackup media storage associated servers are used instead of CloudPoint associated media servers. The NetBackup media storage associated servers must be connected to the CloudPoint server to facilitate all the CloudPoint related operations.

The following table describes the underlying tasks.

Table 3-2 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 22.
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 22.
Discover assets on CloudPoint server	You can discover assets on the CloudPoint server. See “Discover assets on CloudPoint server” on page 27.
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 26.

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 the certificate, key and trust store 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, at `/cloudpoint/openssl/netbackup/`, create the following entries:
 - `ECA_TRUST_STORE_PATH = /cloudpoint/eca/trusted/cacerts.pem`
 - (Optional) `VIRTUALIZATION_CRL_CHECK = CHAIN`
 - (Optional) `ECA_CRL_PATH = /cloudpoint/eca/crl/`

Note: The CA certificates and CRLs should be present under `/cloudpoint/eca/trusted/cacerts.pem` for trust-store, and `/cloudpoint/eca/crl` for CRL.

- 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 DER, PEM, and P7B formats.
- `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.

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, CloudPoint might not work.

- 5 Add the CloudPoint server to NetBackup or run the `tpconfig` command to update the certificate for a CloudPoint server already added to NetBackup.

Add a CloudPoint server

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

Note: To allow backups from snapshots, bi-directional connectivity is required between CloudPoint and NetBackup servers

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 that have provided at the time of CloudPoint installation.
- 9 Click **Save**.

Note: If NetBackup security level is set to VERY HIGH, an additional field **Token** is shown where you can provide a Standard Host Token. This is required for NetBackup certificates generation on CloudPoint. You may need to contact the security administrator or a backup administrator for requesting the additional security permissions required for generating the token.

Add a cloud provider for a CloudPoint server

You can protect the assets on the Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Microsoft Azure Stack Hub 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 the **Providers** tab or click **Add** under the cloud provider for which you want to add a configuration.
- 3** Enter a value in the **Configuration Name** field, in the **Add configuration** pane.
- 4** Select the preferred **CloudPoint server**.

5 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.
	<i>Using AAD:</i> Azure Stack Hub Resource Manager endpoint URL	The endpoint URL in the following format, that allows CloudPoint to connect with your Azure resources. <code>https://management.<location>.<FQDN></code>
	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.
	Authentication Resource URL (optional)	The URL where the authentication token is sent to.
	<i>Using ADFS:</i> Azure Stack Hub Resource Manager endpoint URL	The endpoint URL in the following format, that allows CloudPoint to connect with your Azure resources. <code>https://management.<location>.<FQDN></code>
	Tenant ID	The ID of the AAD directory in which you created the application.

Cloud provider	Parameter	Description
Amazon AWS Note: If the CloudPoint server is configured with IAM Config, the Access Key and Secret Key options are not available.	Client ID	The application ID.
	Secret Key	The secret key of the application.
	Authentication Resource URL (optional)	The URL where the authentication token is sent to.
	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.
Google Cloud Platform	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.
	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.

6 Enter the connection and authentication details in the **Add Configuration** pane.

7 Click **Save**.

The assets on the cloud providers are automatically discovered.

IAM Role for AWS Configuration

If the CloudPoint server 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 22.

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 example, if this is the role ARN for the role you want to use to configure the plugin:

`arn:aws:iam::935923755:role/TEST_IAM_ROLE`

So, to configure the plugin, provide the last part of the ARN, the name: `TEST_IAM_ROLE`

For more details, refer to the *Access 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 **Media server** tab, 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 with a CloudPoint server, automatic discovery is triggered to discover assets from the cloud. During periodic discovery, NetBackup pulls the assets data from CloudPoint every two hours whereas CloudPoint pulls the asset data from cloud provider configurations every one hour. If you disable a CloudPoint server, all the assets associated with that server are no longer protected or synced with NetBackup.

You can also manually trigger the cloud asset discovery if required, using the *Discover* option for individual cloud provider configurations, or you can trigger a discovery on a CloudPoint server to fetch the assets data available on the CloudPoint server.

After the first full discovery, NetBackup subsequently performs periodic incremental discovery of assets for the configured CloudPoint servers. It only detects the changes, such as addition, removal, or modification of assets, that occurred between the last and current discovery.

Note: For the accurate incremental discovery, ensure that the time is set correctly on the NetBackup master server and the CloudPoint server, according to the time-zones they are located in, to avoid any issues with the discovery.

The following procedure describes how to perform discovery at the CloudPoint server level, which does not actually discover the assets from the Cloud, but only fetches the point-in-time data from CloudPoint Server.

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**.

The following procedure describes how to perform discovery at the configuration level, which triggers a deep discovery of assets and fetches the point-in-time state of the assets detecting any additions, modifications, or deletion of assets in the Cloud.

To discover assets for a cloud provider configuration

- 1 On the left, click **Cloud**.
- 2 Click on the **CloudPoint server** tab
- 3 Click the CloudPoint server IP or hostname for which to view the cloud providers.
- 4 Click on the provider tab for which to view the configurations.
- 5 From the menu next to the configuration name, click **Discover**.

Note: If the discovery on cloud provider configurations takes more than 30 minutes, the discovery operation times out. But the subsequent operation continues which syncs the NetBackup assets with the CloudPoint server assets.

Change the autodiscovery frequency for CloudPoint servers

Use `nbgetconfig` and the `nbsetconfig` commands to view, add, or change the autodiscovery option. For example:

`CLOUD_AUTODISCOVERY_INTERVAL = number of seconds`

See the [NetBackup Administrator's Guide, Volume I](#) for more details.

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 In the **Token** field, enter a Reissue token for CloudPoint server
- 6 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**.

Note: After disabling a CloudPoint sever, protection for the associated assets will start failing for that server. In that case, unsubscribe the assets from the protection plans or cancel any pending SLP operations to avoid seeing job failures during the time it is disabled.

(Optional) Add the CloudPoint extension

The CloudPoint extension serves the purpose of scaling the capacity of the CloudPoint host to service a large number of requests concurrently running on the CloudPoint server at its peak performance capacity. You can install one or more CloudPoint extensions on-premise or in cloud, depending on your requirements to run the jobs without putting the host under additional stress. An extension can increase the processing capacity of the CloudPoint host.

The CloudPoint extension can have the configuration same or higher as the CloudPoint host.

Supported CloudPoint extension environments:

- VM based extension for on-premise
- Cloud based extension with managed Kubernetes cluster

Refer to *Deploying CloudPoint extensions* chapter in the latest version of [NetBackup CloudPoint Install and Upgrade Guide](#).

Managing intelligent cloud groups

You can create and protect a dynamic group of assets by defining the intelligent cloud asset groups based on a set of filters called queries. NetBackup selects the

cloud virtual machines, applications, or volumes based on the queries, and adds them to the group. An intelligent group automatically reflects changes in the asset environment and eliminates the need to manually revise the list of assets in the group when the assets are added or removed from the environment.

Then when you apply protection plan to an intelligent cloud asset group, all the assets satisfying the query conditions will automatically be protected if the asset environment changes in future.

Note: You can create, update, or delete the intelligent groups only if your role has the necessary RBAC permissions for the cloud assets that you require to manage. The NetBackup security administrator can grant you access for an asset type (VM, PaaS, application, volume, network) associated with a specific account or subscription, or at a cloud provider level. Refer to the *NetBackup Web UI Administrator's Guide*.

Create an intelligent cloud group

To create an intelligent cloud group

- 1 On the left, click **Cloud**.
- 2 Click the **Intelligent groups** tab and then click **+ Add**.
- 3 Enter a name and description for the group.
- 4 Select the cloud provider, account ID, and region.
- 5 Select the **Asset type**.
- 6 Then do one of the following:
 - Select **Include all assets of the selected type**.
This option uses a default query to select all assets for backup when the protection plan runs.
 - To select only the assets that meet specific conditions, create your own query: Click **Add condition**.

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

See [the section called “Query options for creating intelligent cloud groups”](#) on page 32.

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:

The screenshot shows the NetBackup Intelligent Cloud Groups query builder interface. At the top, 'Asset type' is set to 'Virtual machine'. Below it, there's a checkbox for 'Include all assets of the selected type'. The main query area shows a list of conditions connected by 'AND' and 'OR' operators. The first condition is 'displayName' 'Contains' 'CP'. The second condition is 'tagname' 'Starts with' 'eng'. A third condition is partially visible: 'state' 'is' 'running'. Buttons for '+ Condition', '+ Sub-query', 'Cancel', 'Add and Protect', and 'Add' are visible at the bottom.

This example uses **AND** to narrow the scope of the query: It selects only the VMs that have `cp` in their display name and that also have a tag name as `eng`, and are in `running state`.

Note: Special character '`<`' is not supported in a tag name. If present, asset group creation will fail.

Note: Known limitation in NetBackup - if you create a query that has the asset tag names (referenced from your cloud provider) containing spaces or special characters such as `(,), &, \, /, ", [,], {, }`, you cannot later edit the query for editing any parameters. This does not prevent you from successfully creating the intelligent group and applying the protection plan to it. Only the Edit query functionality is affected with this limitation.

To avoid this issue, ensure that the tag names do not contain the specified special characters and create a new query with the new tag names.

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

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

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

Note: When using queries in **Intelligent groups**, the NetBackup web UI might not display an accurate list of assets that match the query if the query condition has non-English characters.

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

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

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

To save the group and apply a protection plan to it, click **Add and protect**. Select the plan, and click **Protect**.

Query options for creating intelligent cloud groups

Note: The attribute values may not match exactly with values shown on the cloud provider's portal. You can refer to the asset details page or the cloud provider's API response of an individual asset.

Table 3-3 Query keywords

Keyword	Description
	(all values are case-sensitive)
<code>displayName</code>	Asset's display name.
<code>state</code>	For example, running, stopped etc.
<code>tag</code>	A label assigned to the asset for categorization.

Table 3-3 Query keywords (*continued*)

Keyword	Description
	(all values are case-sensitive)
<code>instanceType / machineType / vmSize</code>	Asset's instance/machine type or VM size, depending on the cloud provider selection. For example, t2.large, t3.large, or b2ms, d2sv3

Table 3-4 Query operators

Operator	Description
Starts with	Matches the value when it occurs at the start of a string.
Ends with	Matches the value when it occurs at the end of a string.
Contains	Matches the value you enter wherever that value occurs in the string.
=	Matches only the value that you enter.
!=	Matches any value that is not equal to the value that you enter.

Note: Once you create an intelligent group, you cannot edit the cloud provider selection for it, but you can edit the name and description, and modify the query as required.

Delete an intelligent cloud group

To delete an intelligent cloud group

- 1 On the left, click **Cloud**.
- 2 Locate the group under the **Intelligent 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 **Remove protection**.
- 5 Then select that group under the **Intelligent groups** tab and click **Delete**.

Protecting cloud assets or intelligent cloud groups

You can create the cloud provider-specific protection plans for your cloud workloads. Then you can subscribe the assets associated with that cloud provider to a provider-specific protection plan.

Note: If you previously had a protection plan that was applied to assets from different cloud providers, it will be automatically converted to the new provider-specific format after upgrading to NetBackup 9.1. For example, if you had the assets from Google Cloud and AWS Cloud subscribed to one protection plan, then the protection plan will be split and converted into two separate plans for each provider.

See [the section called “Conversion of protection plans after upgrading to NetBackup 9.1”](#) on page 35. section.

Use the following procedure to subscribe a cloud VM, application, volume, or an intelligent group to a protection plan. When you subscribe an asset to a protection plan, you assign predefined backup settings to the asset.

Note: The RBAC role that is assigned to you must give you access to the assets that you want to manage and to the protection plans that you want to use.

To protect a cloud asset or an intelligent group

- 1 On the left, click **Cloud**.
- 2 On the **Virtual machines** tab, or **Applications** tab, or **Volumes** tab or **Intelligent groups** tab, click the box for the asset or the asset group and click **Add protection**.
- 3 Select a protection plan and click **Next**.
- 4 You can adjust the following settings:
 - **Schedules and retention**
 - **Storage options**
 - **Backup options**
- 5 Click **Protect**.

'Backup now' option for immediate protection

Apart from the scheduled protection plans, you can also use the **Backup now** option to backup an asset immediately, to safeguard against any unplanned circumstances.

1. Select a cloud asset or an intelligent group and click **Backup now**.
2. Then select a protection plan to apply. Only the protection plans relevant to a specific cloud provider of the asset are displayed as options.
3. Click **Start backup**.

A backup job is triggered, which can be tracked on the **Activity monitor** page.

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

Conversion of protection plans after upgrading to NetBackup 9.1

Note the following points with respect to the automatic conversion of older protection plans to the new format.

- Protection plan conversion starts when the assets migration is completed after upgrading NetBackup to 9.1.
- Old protection plans that had no assets subscribed will not be converted to the new format. You can manually delete them.
- **Before or during conversion**
 - All the assets are unsubscribed from the old protection plan and subscribed to the converted protection plan.
 - No new assets can be subscribed to the old protection plan.
 - '*Backup now*' operation will fail for the old plan.
 - Customizing or editing the old protection plan is prevented.
- **After successful conversion**
 - If the old protection plan was used to protect the assets from only one cloud provider, then the new plan retains the same name and asset subscription upon conversion.
 - If the old protection plan was used to protect the assets from multiple cloud providers, then the name of the old protection plan is retained as before, except that it is updated to retain the asset subscription for any one cloud provider upon conversion.
For the other cloud providers which were part of the old plan, new protection plans are created upon conversion, and only the assets of respective providers are subscribed to them. New plans are named in the following format `<old_plan_name>_<cloud_provider>`.
 - Hence you may see more number of plans in your *Protection Plans* menu on the Web UI than before.
 - Success messages are shown in the notifications as follows:

'The protection plan <protectionPlanName> created during conversion to new format.'

'Successfully converted the protection plan <protectionPlanName> to the new format.'

Then you can start managing and applying the converted protection plans as normal.

Failure scenarios

Refer to the following to know how the failure scenarios are handled during or after the conversion of protection plans. Also check the notifications for any failure alerts and take the necessary action.

- Some of the assets might fail to get unsubscribed from the old protection plan. In that case, the conversion still continues with the assets that are successfully unsubscribed. Then the conversion for the assets which failed to get unsubscribed will be retried every four hours.
- After the conversion, some of the assets might fail to get automatically re-subscribed to the new plan. In that case, you need to manually subscribe those assets to the converted protection plan.
- Failure might be encountered while assigning the required access permissions to the new, converted protection plan. In that case, you need to manually assign the access permissions.

Customize or edit protection for cloud assets or intelligent groups

You can edit certain settings for a protection plan, including schedule backup windows and other options.

To customize or edit the protection plan for a cloud asset

- 1** On the left, click **Workloads > Cloud**.
- 2** On the **Virtual machines** tab, or **Applications** tab, or **Volumes** tab or **Intelligent groups** tab, click on the asset that you want to customize the protection for.
- 3** Click **Customize protection > Continue**.
- 4** You can adjust one or more of the following settings:
 - **Schedules and retention**
Change the backup start window.
 - **Backup options**
Enable/disable regional snapshots for Google Cloud assets, or specify/change snapshot destination resource group for Azure and Azure Stack Hub assets.

Remove protection from cloud assets or intelligent groups

You can unsubscribe a cloud asset from a protection plan. When the asset is unsubscribed, backups are no longer performed.

To remove protection from a cloud asset

- 1 On the left, click **Cloud**.
- 2 On the **Virtual machines** tab, or **Applications** tab, or **Volumes** tab or **Intelligent groups** tab, click on the asset that you want to remove the protection for.
- 3 Click **Remove protection > Yes**.

Cloud asset cleanup

Cloud assets are cleaned up automatically during cleanup cycle or manually based on the following criteria:

- No active protection plan to cloud asset.
- Asset is not discovered in last 30 days (cleanup age).
- No recovery points exist.
- Asset is marked for deletion (asset is deleted on CloudPoint server).

User can enhance this cloud asset cleanup criteria by updating cleanup-age and providing specific filter criteria for assets through `bp.conf` file. Following parameters must be configured in `bp.conf` file:

- `CLOUD.CLEANUP_AGE_MINUTES`
- `CLOUD.CLEANUP_FILTER`

For example,

```
/usr/opensv/netbackup/bin/nbsetconfig  
  
nbsetconfig> CLOUD.CLEANUP_AGE_MINUTES = 180  
  
nbsetconfig> CLOUD.CLEANUP_FILTER = "provider eq 'aws'"  
  
nbsetconfig>
```

User can also manually run POST named query cleanup-assets with the following request body and then run GET with query-id from POST response as described in the following example:

```
{  
  "data":{
```

```

    "type": "query",
    "attributes": {
      "queryName": "cleanup-assets",
      "workloads": ["cloud"],
      "parameters": {
        "cleanup_age_minutes": 180
      },
      "filter": "provider eq 'aws'"
    }
  }
}

```

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	GovCloud (US) regions
Amazon Web Services	<ul style="list-style-type: none"> ■ us-gov-east-1 ■ us-gov-west-1
Microsoft Azure	<ul style="list-style-type: none"> ■ 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 22.

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 the following:

<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.

Note: You can now directly associate a snapshot to an existing peer resource group, at the time of creating a protection plan. However the functionality of defining a peer resource group by specifying a prefix which is described in this section, still exists.

Refer to information on creating protection plans in the *NetBackup Web UI Administrator's Guide* for the complete procedure.

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 3-5 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	<p>NetBackup creates new the 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 the conditions are not met, snapshot jobs fail.</p>

Table 3-5 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 3-6 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. 	<p>Snapshots are created in the peer resource group.</p>

Table 3-6 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 3-6 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. 	Snapshots are not created and the job fails.
<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. 	Snapshots are not created and the job fails.

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:

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

About the NetBackup Accelerator for cloud workloads

NetBackup Accelerator reduces the backup time for cloud backups. NetBackup uses reference snapshots 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.

NetBackup supports Accelerator backup for AWS, Azure and Azure Stack workloads.

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).
- Reduces the I/O on the CloudPoint server.
- Reduces the CPU load on the CloudPoint server.

How the NetBackup Accelerator works with virtual machines

For Azure and Azure Stack backups, Accelerator is activated when you select a Accelerator supported storage type, like MSDP, OpenStorage, CloudStorage, and MSDP-C (Azure and AWS).

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.
- 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 VM backup. The changed blocks are identified by comparing the previous reference snapshot and the current snapshot. If you select **Keep backup only** or **Initiate backup when snapshot is about to expire**

option in the protection plan, the snapshot is retained for accelerator purpose till the next backup is completed.

- 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.
- Microsoft Azure does not allow more than 200 subsequent incremental snapshots. If you select the **Keep snapshot along with backup** option in the protection plan and specify a such a retention period for the snapshot, so that it leads to more than 200 incremental snapshots. Then, full backups take place instead of accelerator. It is recommended to keep a reasonable snapshot retention period to utilize the accelerator benefits.
- If the configuration of a VM changes, for example, if a new disk is added to a VM between two accelerator backups, a full backup is taken for that disk, and accelerator backup is taken for the existing disks.

Accelerator forced rescan for virtual machines (schedule attribute)

Accelerator forced rescan helps to prevent corrupt backup image issues by manually executing the ForcedRescan command. 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.

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.

Recommendations for using forced rescan:

- Do not trigger force rescan for the VMs which are turned off.
- If the storage location memory is full, you can see a notification in the UI. Initiate the force rescan only when sufficient memory is available at the storage location.

NetBackup creates a schedule named 'ForcedRescan' for every protected VM. To manually trigger the backup with force rescan execute the following command in the command prompt or the Linux terminal:

```
bpbackup -i -p <policy_name> -s ForcedRescan
```

For example, `bpbackup -i -p`

```
msdp_10mins_FRS+5d990ab5-f791-474f-885a-ae0c30f31c98 -s ForcedRescan
```

You can obtain the policy name from web UI from the relevant protection plan.

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.

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:

```
Jul 21, 2021 1:55:52 PM - Info bpbrm (pid=78332) accelerator enabled
Jul 21, 2021 1:55:53 PM - Info bpbrm (pid=78332) There is no
complete backup image match with track journal, a regular full
backup will be performed.
```

..

```
Jul 21, 2021 1:56:11 PM - Info bpbkar (pid=1301) accelerator sent
402666496 bytes out of 402664960 bytes to server, optimization 0.0%
```

When subsequent backups of the virtual machine use Accelerator, the following messages appear in the job details log:

```
Jul 21, 2021 2:01:33 PM - Info bpbrm (pid=79788) accelerator enabled
```

..

```
Jul 21, 2021 2:02:00 PM - Info bpbkar (pid=1350) accelerator
sent 1196032 bytes out of 402664960 bytes to server, optimization 99.7%
```

This message is a key trace for Accelerator. In this example Accelerator was successful at reducing the backup data by 99.7%.

Configuring backup schedule for cloud workloads

You can add backup schedule in the Attributes tab of the Add backup schedule dialog, while creating a protection plan for the Azure, Azure Stack, and AWS cloud workloads.

See the *Managing protection plans* section of the *NetBackup Web UI Administrator's Guide*, for details of how to create a protection plan.

To add backup schedule to a cloud workload

- 1 On the left, click **Protection > Protection plans** and then click **Add**.
- 2 In **Basic properties**, enter a **Name**, **Description**, and select **Cloud**, from the **Workload** drop-down list.
- 3 Select a **Cloud Provider** from the drop-down list, click **Next**. In **Schedules**, click **Add schedule**.

In the **Add backup schedule** tab, you can configure the options for retaining the backup and the snapshot.

- 4 From the **Recurrence** drop-down, specify the frequency of the backup.
- 5 In the Snapshot and backup options, do any of the following:
 - Select **Keep snapshot along with backup** option to retain both the snapshot and the backup. Specify retention period for both the snapshot and the backup, using the **Keep snapshot for** and the **Keep backup for** drop-downs. Select **Full** from the **Backup type** drop-down. Select **Initiate backup only when the snapshot is about to expire** option, to start the backup job just before the retained snapshot expires.
 - Select **Keep snapshot only** option, to retain only the snapshot. Specify retention period for the snapshot using the **Keep snapshot for** drop-down.
 - (Optional) If you have selected provider as Amazon AWS, and selected to retain the snapshot by selecting any of the above two options, you can configure snapshot replication at this point. For more information about cloud snapshot replication, See [“Configure AWS snapshot replication”](#) on page 52.
 - Select **Enable Snapshot replication**.
 - In the table, select **Region**, **AWS Account**, and **Retention** period for the replicated snapshots.

Note: The number of replication copies that you configure is displayed in the **Snapshot replicas** column in the **Schedules and retention** table in the **Schedules** tab.

- Select **Keep backup only** option, to retain only the backup. The snapshot expires immediately after the backup. Specify retention period for the backup using the **Keep backup for** drop-down. Select **Full** from the **Backup type** drop-down.

Note: As NetBackup supports granular restore only from the snapshot, if you select **Keep Backup Only** the granular recovery options do not work. Similarly, AWS snapshot replication feature does not work if you select **Keep Backup Only**.

- 6 Continue creating the schedule in the **Start window** tab, as described in the *Managing protection plans* section of the *NetBackup Web UI Administrator's Guide*.

Availability of granular recovery for different backup options

Availability of the granular recovery for files or folders option, depends on the different backup options that you select for the workload.

- When you select the **Keep snapshot along with backup option**, granular recovery is available.
- When you select the **Keep snapshot only** option, granular recovery is available.
- When you select the **Keep backup only** option, granular recovery is not available.

Indexing during backup and snapshot jobs

- NetBackup performs VxMS (Veritas Mapping Service) based indexing from snapshot, and inline indexing during the backup from snapshot Jobs. It can index files irrespective of the region and location of the CloudPoint server. VxMS based indexing is currently supported for AWS, Azure, and Azure Stack Hub clouds.
- Indexing is performed during the actual backup or snapshot jobs, but you can perform the recovery of individual files or folders only from the snapshot copy using **Enable granular recovery for files and folders** option.
- Once the snapshot of the VM assets is created, the 'Index from Snapshot' job for each of the assets is triggered. You can check the indexing job details in the **Activity Monitor**.
- The VxMS debug logs and the cloud connector debug logs are available in the `/cloudpoint/openv/dm/datamover.<datamover-id>/netbackup/logs` folder of the CloudPoint server.

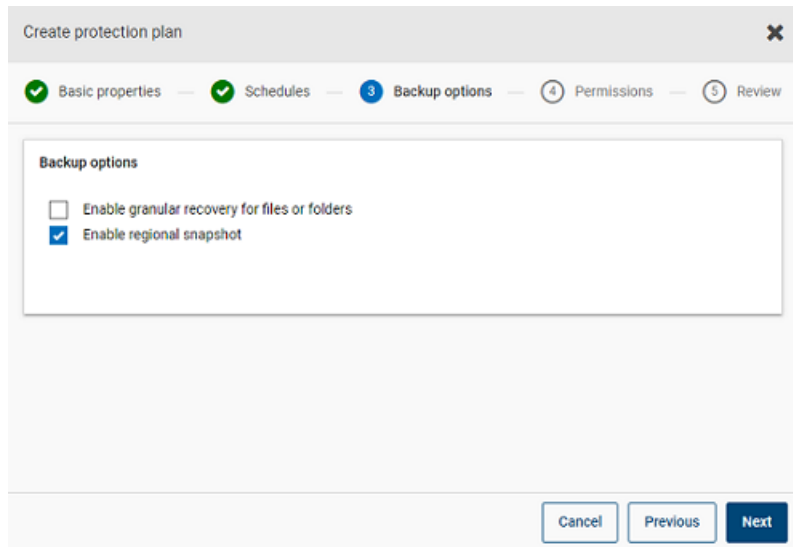
Note: If the VM is not in connected state, then the VM backup continues and the backup job is marked as partially successful. In this case, you cannot restore individual files or folders as the indexing is not available when the VM is not connected.

Backup options for cloud workloads

Regional snapshots for Google cloud

You can choose to enable regional snapshots for the Google cloud workloads while creating a protection plan.

If the regional snapshot option is enabled, the snapshot will be created in the same region in which the asset exists. Otherwise, the snapshot will be created in a multi-regional location.



Create protection plan

Basic properties — Schedules — 3 Backup options — 4 Permissions — 5 Review

Backup options

- ☐ Enable granular recovery for files or folders
- ☒ Enable regional snapshot

Cancel Previous Next

Snapshot destination resource group for Azure and Azure Stack Hub

You can choose to specify a snapshot destination peer resource group while creating a protection plan for Azure or Azure Stack Hub. While the previous functionality of defining a peer resource group by specifying a prefix still exists, you can now directly associate a snapshot to an existing peer resource group at the time of creating a protection plan.

If you have selected the cloud provider as Microsoft Azure or Azure Stack Hub while creating a protection plan, you can select **Specify snapshot destination resource group** to associate snapshots to a particular peer resource group within the same region in which the asset exists. Then select a configuration, subscription, and a resource group for a snapshot destination.

The snapshot is stored in one of the destination resource groups, in the following preference:

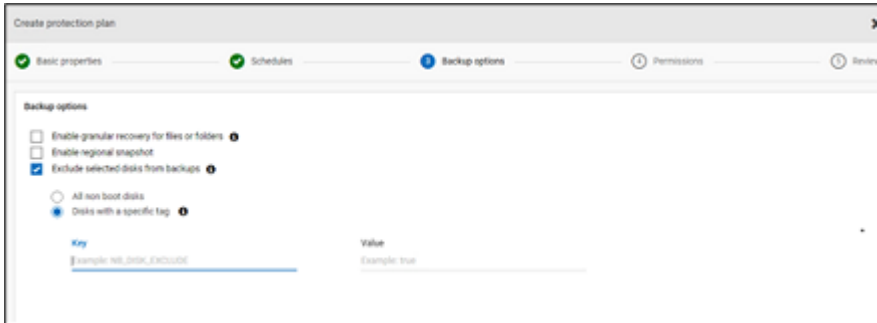
- A destination resource group specified in the protection plan
- A pre-fixed resource group specified in the plugin configuration (for Azure only)
- A resource group in which the asset exists, if no destination or pre-fixed resource group is specified in NetBackup.

Excluding selected disks from backup

You can configure a protection plan to exclude some disks from the backup and snapshot which are applicable to all supported cloud vendors including GCP. This enables you to avoid redundant images of the disks that do not need to be backed up, and speed up the backups by reducing the volume of data to be processed.

If you are creating a protection plan for AWS, Azure, Azure Stack Hub, or GCP clouds, you can select **Exclude selected disks from backups** option and specify the disks that should not be included in the backup image. You can choose to

exclude either all the non-boot disks, or the disks that have specific tags associated with them in the corresponding cloud provider account.



Note: A protection plan that has disk exclusion option enabled can be applied only to the cloud VM type assets and VM intelligent groups.

Then while restoring the VMs from the Recovery Points tab, refer to the **Includes disks** column to view the list of disks that are included or excluded in the backup image.

Refer to the information on creating a protection plan in the *NetBackup Web UI Administrator's Guide* for the complete procedure.

Notes:

- In case of LVM, if disks are excluded partially then system might not boot up properly.
- In case there is a non supported file system configured on a disk and user wants to exclude that disk from snapshot, the snapshot would continue to be crash consistent snapshot even if the disk containing non supported file system is excluded.
- If the user wants to exclude the disk, he/she should have the **nofail** flag attached to the data disk prior to taking a snapshot in the `/etc/fstab` file. This is required if the user reboot the instance without this volume attached (for example, after moving the volume to another instance), the **nofail** mount option enables the instance to boot even if there are errors mounting the volume. For more information, refer to the following example entry in the `/etc/fstab` file:
For example, **UUID=aebf131c-6957-451e-8d34-ec978d9581ae /data xfs defaults,nofail 0 2**
- The user should ensure that assets are properly discovered once there has been any change in their tags from the cloud provider. Once the policy run is

scheduled for an asset, the disks are excluded as per the discovered data only. If, the user attaches a tag while the snapshot is in progress, that **tag** would not be considered as a part of exclusion. Once discovery is completed, it will be considered during next protection cycle.

- In case of OS with non-English locale, if user opts tag based exclusion in protection plan and disk tag has non-English character, even then disk exclusion works as expected. But in some cases tag with non-English character is not correctly captured in job(try) logs and audit logs although there is no functionality impact as disk exclusion is considered correctly.

Snapshot replication

Replicating a snapshot means saving a copy of the snapshot to another location. In AWS, another location can be one of the following:

- different region within the same account.
- same region in a different account.
- different region within different account.

For example, an AWS cloud administrator have their assets in the region X. The snapshots of those assets will also be stored in X region. However, you can also replicate the snapshots to the Y region within same account or X/Y region in a different account, for an added level of protection. In NBU CloudPoint terminology, the original location (X) is the replication source, and the location where snapshots are replicated (Y) is the replication destination.

Replication is performed in three steps. This mechanism is handled internally and the entire process is completely transparent to the user.

- Share the snapshot, only if replicating to a cross account. For more information, see the [Share a snapshot](#) section of the AWS documentation.
- Copy the snapshot. For more information, see the [CopySnapshot](#) section of the AWS documentation.
- Unshare the snapshot, only if replicating to a cross account.

Configure AWS snapshot replication

Requirements for replicating snapshots

- Replicating unencrypted snapshots

Ensure that the source and target accounts/regions are configured using the AWS cloud provider from NetBackup CloudPoint. There are no additional requirements for replicating unencrypted snapshots.

- **Replicating encrypted snapshots using AWS KMS**

Ensure that the source and target accounts/regions are configured using the AWS cloud provider from NetBackup CloudPoint.

Additionally, to replicate encrypted snapshots to a cross account, the encryption CMK key from the original location needs to be shared to the target account. (This shared KMS key is implicitly used while copying the snapshot in the target account, and the copied snapshot can be replicated by a different key).

Both the source and target locations should have encryption key (KMS key) with same name; that is, they should have the same key alias (in terms of AWS).

If encryption key with the same name is not present at the target, then the replicated snapshot is encrypted using the default KMS key in the target location.

- **Permissions for cross account replication**

For cross-account replication, the AWS IAM user or role associated with the snapshot source region's AWS account (source AWS account) must have the following permissions:

- `ModifySnapshotAttribute` and `CopySnapshot` on the EC2 instance.
- `DescribeKey` and `ReEncrypt` on the KMS key that is used to encrypt the original snapshot.

For cross-account replication, the AWS IAM user or role associated with the snapshot replication target region's AWS account (target AWS account) must have the following permissions:

- `CreateGrant`, `DescribeKey`, and `Decrypt` on the KMS key that is used to encrypt the original snapshot.
- `CreateGrant`, `Encrypt`, `Decrypt`, `DescribeKey`, and `GenerateDataKeyWithoutPlainText` on the KMS encryption key used while performing the `CopySnapshot` operation on the original snapshot.

You can choose to replicate snapshots for AWS cloud assets from the primary location to a remote or a secondary location. The CloudPoint servers 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 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:

- 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 CloudPoint server must be replicated to the region that is discovered by the same CloudPoint server.
For example, assets that are discovered by CloudPoint server 'CP1' cannot be subscribed to a protection plan replicating to a region that is discovered by CloudPoint 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.

Using AWS snapshot replication

This section elaborates how to create snapshot replicas using the AWS snapshot replication feature, and restore the replicated snapshots whenever required. Refer to the *NetBackup™ CloudPoint Install and Upgrade Guide* and the *NetBackup Web UI Administrator's Guide* for details about these steps, otherwise indicated.

Creating snapshot replications

This section describes how to configure the Source region to create snapshot replicas in the Target region.

To create replicas

- 1 Add CloudPoint server (CP1) in webUI. See [“Add a CloudPoint server”](#) on page 22.
- 2 Add AWS plug-in for Source and Target region for replication.
- 3 Create protection plan and select **Region** and **Account**. See [“Configuring backup schedule for cloud workloads”](#) on page 47.
- 4 Connect and configure an application consistent guest VM using the OnHost agent.
- 5 Trigger snapshot-based backup and replicate the snapshots using the protection plan.
- 6 Verify the recovery points for snapshot and replica copy.

Restoring from the snapshot replicas in the target region

If the Source region fails, you can restore the VMs belonging to this region, from the Target region, where you have taken the snapshot replicas. As the Source region is down, you will initially need to restore the VMs in the Target region.

Note: You cannot restore single files or folders from a replica that was discovered by an alternate CloudPoint server in a failed over region.

Restoring in the target region

- 1 Disable server CP1 in the Source region from webUI. See [“Enable or disable a CloudPoint server”](#) on page 29.
- 2 Register a new CloudPoint server (CP2) in the target region, from webUI.
- 3 Add AWS plug-in for only the Target region and account. Let the discovery complete.
- 4 To restore VMs:
 - Log on to the NetBackup console.
 - On the left, click **Cloud**, under **Workloads**. On the **Virtual machines** tab, click the machine that you want to recover.
 - Click the **Recovery points** tab. In the list of images, click **Restore** in front of the required **Replica** image, and click **Restore virtual machine**.
 - To change the Display name for the VM, enter a new name.
 - Select a subnet (subnet path having VPC).
See [“Recovering cloud assets”](#) on page 62.
- 5 Add appropriate security group to the restored VMs to enable remote access.
- 6 Uninstall and reinstall the CloudPoint agent from the restored VMs, and then register the CloudPoint agents with the new CP2 server.
- 7 Run a deep discovery from the AWS provider console.
- 8 Create new protection plan to protect the restored VMs. Trigger a snapshot-based backup.

Restoring back to the source region from the target region

You can restore the VMs from the Target region to the Source region, once the source region is back online.

Restoring to the source region

- 1 Edit the AWS plug-in for CP2 and add the Source region.
- 2 Create a new protection plan to create a snapshot replica in the Source region.
- 3 Trigger snapshot-based backup and replicate.
- 4 Disable the CP2 server in webUI. See [“Enable or disable a CloudPoint server”](#) on page 29.

- 5
- Enable the CP1 server and trigger deep discovery from AWS provider console.
- 6
- Perform full restore of the VMs from the Target region.
- 7
- Add appropriate security group to enable remote access to restored VMs.
- 8
- Uninstall and reinstall the CloudPoint agents from the restored VMs, and then register CloudPoint agents with the CP1 server.
- 9
- Run a deep discovery from the AWS console.
- 10
- Use the existing protection plan to protect newly restored VMs.

Support matrix for account replication

Table 3-7 Support matrix for same account replication

Asset types	Source asset (Region X)	Source snapshot (Region X)	Replicated snapshot (Region Y)
EBS Volume, EC2 Instance and RDS/Aurora	Unencrypted	Unencrypted	Unencrypted
	Attached disks encrypted using default AWS KMS key.	Attached disks encrypted using default AWS KMS key.	Attached disks encrypted using default AWS KMS key.
	Encrypted using AWS KMS CMK key (with Alias ABC).	Encrypted using AWS KMS CMK key (Alias ABC).	Encrypted using AWS KMS CMK key with name if present (Alias ABC), else encrypted using default AWS KMS key.

Table 3-8 Support matrix for different account same region replication

Asset types	Source asset (Account A Region X)	Source snapshot (Account A Region X)	Replicated snapshot (Account B Region Y)
EBS Volume, EC2 Instance and RDS/Aurora	Unencrypted	Unencrypted	Unencrypted
	Encrypted using default AWS KMS key.	Encrypted using default AWS KMS key.	Not supported
	Encrypted using AWS KMS CMK key (with Alias ABC).	Encrypted using AWS KMS CMK key (with Alias ABC).	Encrypted using AWS KMS CMK key with name if present (with Alias ABC), else encrypted using default AWS KMS key.

Table 3-9 Support matrix for different account different region replication

Asset types	Source asset (Account A Region X)	Source snapshot (Account A Region X)	Replicated snapshot (Account B Region Y)
EBS Volume and EC2 Instance	Unencrypted	Unencrypted	Unencrypted
	Encrypted using default AWS KMS key.	Encrypted using default AWS KMS key.	Not supported
	Encrypted using AWS KMS CMK key (with Alias ABC).	Encrypted using AWS KMS CMK key (with Alias ABC).	Encrypted using AWS KMS CMK key with name if present (with Alias ABC), else encrypted using default AWS KMS key.

Table 3-9 Support matrix for different account different region replication
(continued)

Asset types	Source asset (Account A Region X)	Source snapshot (Account A Region X)	Replicated snapshot (Account B Region Y)
RDS	Unencrypted	Unencrypted	Unencrypted
	Encrypted using default AWS KMS key.	Encrypted using default AWS KMS key.	Not supported
	Encrypted using default AWS KMS key.	Encrypted using default AWS KMS key.	Not supported
Aurora	Unencrypted	Unencrypted	Not supported
	Encrypted using default AWS KMS key.	Encrypted using default AWS KMS key.	Not supported
	Encrypted using default AWS KMS key.	Encrypted using default AWS KMS key.	Not supported

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 alternate location restore of MongoDB and MS SQL workloads, the target host must be discovered but the application status should not be in connected or configured state.
- For alternate location restore of Oracle workloads, the target host must be discovered but the application status should not be in connected or configured state.

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 **Manage credentials**.
 - Enter the credentials. If the credentials for the VM are not configured, you must configure the credentials. See the *Managing credentials* chapter of the *WebUI Administrator Guide*.
 - After the virtual machines are connected, the virtual machines state is updated to **Connected**.
- 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 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 **Manage credentials**.
- 3 Update the credentials.

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**.

Discovering PaaS assets

NetBackup allows you to discover and restore Azure SQL database and Azure SQL managed database assets backed up by Microsoft Azure. The supported backup modes are Point in time backup and Long-term retention backup.

Note: You cannot add PaaS assets to any NetBackup protection plans.

To discover PaaS assets

- 1 Add a CloudPoint server. See [“Add a CloudPoint server”](#) on page 22.
- 2 Add Microsoft Azure as a provider. See [“Add a cloud provider for a CloudPoint server”](#) on page 22.
- 3 Run a discovery. See [“Discover assets on CloudPoint server”](#) on page 27.

After the discovery is complete, you can find all the discovered Azure SQL database and Azure SQL managed database assets in the **PaaS** tab, in the **Cloud** workload.

Note: When you create and delete a PaaS assets with same name in intervals, and if the PaaS asset is deleted after discovery, webUI shows old data until the next periodic discovery runs.

Recovering cloud assets

This chapter includes the following topics:

- [Recovering cloud assets](#)
- [Perform rollback recovery of cloud assets](#)
- [Recovering PaaS assets](#)

Recovering cloud assets

You can restore AWS, Azure, and Azure Stack VM assets from snapshot copy, replica copy, backup copy, or duplicate copy.

While restoring VMs, NetBackup gives you the option to change certain parameters of the original backup or snapshot copy. Including options like changing the VM display name, changing power options of the VM, removing tag associations during restore, and restoring to an alternate network. You can also restore VMs to an alternate configuration, to a different region, to a different subscription, and restore VMs or disks to a different resource group.

About the pre-recovery check for VMs

Pre-recovery check indicates how a restore may fail, before the restore is initiated. The pre-recovery check verifies the following:

- Usage of supported characters and the length in the display name.
- Existence of destination network
- Existence of selected Resource group for VMs and disks
- Existence of source VM snapshot(applicable for restore from snapshot)
- Existence of the staging location added in the file
`/cloudpoint/azurestack.conf` (applicable for restore from backup for Azure stack)

- Existence of a VM with the same display name.
- Connectivity with the CloudPoint server and cloud credential validation.

About the pre-recovery check for PaaS assets

The pre-recovery check for PaaS assets verifies the following:

- A valid Azure SQL database display name. Another database with the same display name should not exist.
- Valid restore point in Azure. The restore point must be later than or equal to earliest recovery point and earlier than or equal to the current time.
- A valid managed instance in Azure. The specified managed instance must exist. If the managed instance exist then only SQL database display name is validated under it. This is applicable only for Azure SQL Managed database recovery points.

Supported parameters for restoring cloud assets

The following table summarizes the different parameters that you can change while restoring assets for different cloud providers.

Table 4-1 Supported parameters for Azure and Azure Stack snapshot and backup copies

Parameters	Snapshot copy		Backup copy		
	Azure	Azure Stack	Azure	Azure Stack	AWS
Change VM display name	Y	Y	Y	Y	Y
Change power state of the VM	Y	Y	Y	Y	Y
Remove tag associations	Y	Y	Y	Y	Y
Restore to a different network	Y	Y	Y	Y	Y
Subscription ID			Y	Y	

Table 4-1 Supported parameters for Azure and Azure Stack snapshot and backup copies (*continued*)

Change resource group	Y	Y	Y	Y	
Change region of the VM			Y	Y	Y
Change provider configuration			Y	Y	Y
Change resource group for disks	Y	Y	Y	Y	

Table 4-2 Parameters supported by AWS and GCP snapshot copies

Parameters	AWS	GCP
Change VM display name	Y	Y
Change power state of the VM	Y	Y
Remove tag associations	Y	Y
Restore to a different network	Y	Y

Recovering virtual machines

To recover a VM

- 1 On the left, click **Cloud**.
- 2 Click the **Virtual Machines** tab.
All the discovered cloud assets for the respective category are displayed.
- 3 Double-click 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. For AWS workloads you can see replica as well as backup images, if available.

- 5 In the **Copies** column, click the copy that you want to recover. You can see the backup, snapshot, and replica copy, if available. Click **Recover**. If you don't select a copy to restore, the primary copy is selected.
- 6 Click **Restore Virtual Machine**.
- 7 In the Recovery target page, do the following:
If you restore a backup copy, modify the values of these parameters as required:
 - **Configuration**: To restore to an alternate configuration, select one from the drop-down.
 - **Region**: To restore to an alternate region, select one from the drop-down.
 - **Subscription**: To restore to an alternate subscription, select one from the drop-down. For Azure and Azure Stack only.
 - **Resource group**: To restore to an alternate resource group, click the search icon, in the **Select resource group** dialog, select the required resource group. For Azure and Azure Stack only.
 - **Display name**: To change the display name, enter the new one in the field. The specified display name is validated during the pre-recovery check.

Note: Except in AWS workloads, the following special characters are not allowed in the display name: ` ~ ! @ # \$ % ^ & * () = + _ [] { } \ \ | ; : ' \" , < > / ? ."

If you restore a snapshot copy, specify only the **Resource group** and the **Display name**.

- 8 Click **Next**.
- 9 In the Recovery options page:
 - If you restore a backup copy, to restore to a different region, select a **Region**. To select a network available in that region, click the search icon near **Network configuration**, and select a target network for recovery.
 - If you restore a snapshot copy, click the search icon in **Network configuration**, and select a target network for recovery. The list shows networks available in that region.

In the **Advanced** section:

- To keep the VM powered on after recovery, select **Power on after recovery**.
- To remove the tags associated with the asset at the time of backup or creating snapshot, select **Remove tag associations**.

Note: If you do not select the **Remove tag associations** option, any tag value for assets should not have spaces, before and after a comma. After the restoration of an asset, the spaces before and after a comma in the tag values are removed. For example, the value for the tag name: **created_on:** *Fri, 02-Apr-2021 07:54:59 PM , EDT*, is converted to: *Fri,02-Apr-2021 07:54:59 PM,EDT*. You can manually edit the tag values to reinstate the spaces.

- 10 Click **Next**. The pre-recovery check begins. This stage validates all the recovery parameters and displays errors, if any. You can fix the errors before starting the recovery.

- 11 Click **Start recovery**.

The Restore activity tab shows the job progress.

For information on the recovery status codes, see the NetBackup administrator or the *NetBackup Status Codes Reference Guide*, available here:

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

Recovering applications and volumes to its original location

For GCP, 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 applications and volumes to the original location

- 1 On the left, click **Cloud**.
- 2 Click the **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. 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 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.

Recovering applications and volumes to an alternate location

Considerations

- For encrypted VM restore in AWS 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 applications and volumes to alternate location

- 1 On the left, click **Cloud**.
- 2 Click the **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. 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 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 On the left, click **Activity monitor** to view the job status.

Recovery scenarios for GCP VMs with Read-only volumes

The following table describes how NetBackup handles the restore/recovery of Google Cloud Platform VMs that have Read-only volumes.

Scenario	Handling
Restoring a volume from the snapshot of an attached Read-only disk (from the <i>Volumes</i> tab under Cloud workloads on the NetBackup Web UI)	During restore, the disk will be attached in the 'Read/write' mode to the original or alternate location.
Restoring a VM (with a Read-only disk) from a crash-consistent snapshot (from the <i>Virtual Machines</i> tab under Cloud workloads on the NetBackup Web UI)	During restore of such a VM to its original or alternate location, a 'Read-only' disk will be restored in a 'Read/write' mode.

Scenario

Restoring a VM (with a Read-only disk) from an app-consistent snapshot (from the *Virtual Machines* tab under Cloud workloads on the NetBackup Web UI)

Handling

Even if a 'Read-only' disk can be attached to multiple VMs, it will be discovered under only one VM.

For a Windows VM, the snapshot will fail with a VSS error, similar to the following:

```
Failure: flexsnap.GenericError:
Failed to take snapshot(error:
Failed to create VSS snapshot of
the selected volumes.)"
```

For a Linux VM, the snapshot may or may not be successful for a VM under which the disk is discovered, but it will fail for other VMs due to the missing dependencies. Error example:

```
linear_flow.Flow: create snapshot
(test-win) of host linux-1(len=4) '
requires
['snap google-gcp-us-west2-b-7534340043132122994']
but no other entity produces said
requirements\n MissingDependencies
```

In the above case, if a snapshot is successful for a Linux VM, a 'Read-only' disk will be restored in a 'Read/write' mode.

Perform rollback recovery of cloud assets

The rollback recovery of a cloud asset overwrites the existing data on the original asset. Unlike virtual machine restore, rollback restore does not create a new copy of the restored image, but replaces the existing data on the source.

Note: Snapshot replicas do not support rollback. Also, Azure Stack and GCP workloads does not support rollback restore.

To perform rollback recovery of the cloud asset

- 1 On the left, click **Cloud**.
- 2 Click 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. The available images are listed in rows with the backup timestamp for each image. In the **Copies** column, click the snapshot that you want to recover. Click **Recover** > **Rollback restore**.
- 5 Click **Start recovery**. The existing data is overwritten.
- 6 On the left, click **Activity monitor** > **Jobs** to view the job status.

Recovering PaaS assets

NetBackup allows you to restore Azure SQL database and Azure SQL managed database assets backed up by Microsoft Azure. The supported backup modes are Point in time backup and Long-term retention backup.

Note: Restoration in Elastic pool in Instance pool are not supported.

Before proceeding make sure that you have the required permissions to restore PaaS assets.

To recover point in time backup assets:

- 1 On the left, click **Cloud**.
- 2 Click the **PaaS** tab.

All the discovered PaaS assets are displayed.
- 3 Click **Restore** in the row of the protected asset that you want recover.
- 4 In the **Recovery points** tab, under **Point in time backup**, click **Restore**.
- 5 Select a date and time under **Restore point (UTC)**. You can select any restore point, between the earliest restore point, and the:
 - latest backup time for online databases.
 - database deletion time for deleted databases.

Microsoft Azure may round off the selected time to the nearest available recovery point, using UTC time.

The default restore date and time displayed in webUI, may differ based on the selected PaaS asset. For example, for Azure SQL databases, the default restore time is the current time, and for Azure SQL managed database, the default restore time is 6 minutes earlier than the current time.

- 6 Optionally, for Azure SQL databases, enter a name for the restored database in the **Database name** field. Database names cannot have special characters like <>*&:\ / and ? or control characters. Do not end the name with a period or space. For more information about Azure resource naming rules, see <https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/resource-name-rules#microsoftsql>

If you do not enter a name, NetBackup automatically assigns a name in the <dbName>_<Restored time in UTC> format.

- 7 Optionally, for Azure SQL managed databases, enter the instance name in the **Managed instance** field. By default, the instance name of the recovery point is displayed. You can also search for the managed instance name using the search option. You can restore to the same region to which your subscription belongs.

If you cannot see the desired managed instance in the search results, perform a manual discovery. Also, ensure that you have RBAC access to the managed instance.

- 8 Click **Next**. Once the Pre-recovery check is complete, click **Start recovery**.
You can check the status of the job in the activity monitor.

To recover long term retention backup assets:

- 1 On the left, click **Cloud**.
- 2 Click the **PaaS** tab.
All the discovered PaaS assets are displayed.
- 3 Click **Restore** in the row of the protected asset that you want recover.
- 4 In the **Recovery points** tab, under **Long term retention backup**, click **Restore** against the image that you want to restore.
- 5 Optionally, for Azure SQL databases, enter a name for the restored database in the **Database name** field. Database names cannot have special characters like <>*&:\ / and ? or control characters. Do not end the name with a period or space. For more information about Azure resource naming rules, see <https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/resource-name-rules#microsoftsql>

If you do not enter a name, NetBackup automatically assigns a name in the <dbName>_<Restored time in UTC> format.

- 6 Optionally, for Azure SQL managed databases, enter the instance name in the **Managed instance** field. By default, the instance name of the recovery point is displayed. You can also search for the managed instance name using the search option. You can restore to the same region to which your subscription belongs
- 7 Click **Next**. Once the Pre-recovery check is complete, click **Start recovery**.
You can check the status of the job in the activity monitor.

Note: Tags from portal as well as CloudPoint are not restored. However, the "createdby: cloudpoint" tag is created while restoring through NetBackup.

Note: For provider protected recovery jobs, any intermittent failures keep the recovery job running until the next schedule job cleanup executes.

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 from cloud virtual machines](#)
- [Restoring volumes on cloud virtual machines](#)
- [Performing steps after volume restore containing LVM](#)
- [Troubleshooting](#)

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 5-1 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 5-2 Supported versions

Application	Version
NetBackup	10.0
NetBackup backup host OS	RHEL 7.x and 8

Table 5-2 Supported versions (*continued*)

Application	Version
CloudPoint host OS	<ul style="list-style-type: none"> ■ RHEL 7.x and later, RHEL 8.4 and 8.5 ■ Ubuntu 18.04 LTS and 20.04 LTS <p>Note: The version of the OS (Ubuntu 20.04 LTS) listed on the UI is the version of the container.</p>
Cloud providers	<ul style="list-style-type: none"> ■ Amazon Web Services ■ Microsoft Azure ■ Microsoft Azure Stack Hub ■ Google Cloud Platform
CloudPoint or agent instance type	<ul style="list-style-type: none"> ■ Amazon AWS: t2.large/t3.large ■ Microsoft Azure: D2s_V3Standard ■ Microsoft Azure Stack Hub: DS2_v2 Standard, DS3_v2 Standard ■ Google Cloud Platform: n1.Standard2 and larger
CloudPoint agent host to be protected	<ul style="list-style-type: none"> ■ Linux OS: RHEL 7.x and RHEL 8.2, 8.4 and 8.5 ■ Windows OS Version: 2012 R2, 2016, 2019 and 2022

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) Note: Granular restore for RHEL 8.3 and 8.2 agent host is supported only when CloudPoint is deployed on RHEL 8.3.	<ul style="list-style-type: none"> ■ ext3 ■ ext4 ■ xfs 	<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.

Note: GRT is allowed irrespective of destination file-system/partition type (FAT, ReFS, LDM or LVM).

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 and Azure Stack Hub) Even if CloudPoint is not deployed in the same subscription and region as the connected VM, but if a backup schedule is configured as part of the protection plan, then granular restore can be performed. For snapshot-only protection plan schedule, for both Azure and Azure Stack Hub, you need to deploy the CloudPoint host in the same subscription and region as the VMs.
- 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 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 indexable flag is applicable for a snapshot-only protection plan schedule.
- 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 74.
- Configure port 5671 and 443 for open CloudPoint host.

- For agentless restore, in Linux systems, configure the port 22 on the indexable virtual machines. For Windows platform, configure the ports 135, 445 and the dynamic/fixed WMI-IN port on the indexable virtual machines.
- 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.
- (Applicable for GCP only) 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 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 LDM disks are ignored when host consistent snapshots are created and indexed.
- Until old agent (pre-installed) service is not restarted, alternate host restore (GRT and application) of LVM asset might fail. You need to restart the older agents in order to support recovery of LVM assets.
- Granular restore (GRT) or single file restore (SFR) can be performed with the help of VxMS indexing. VxMS indexing is applicable for all CloudPoint supported file systems. VxMS indexing can be performed for Azure, AzureStack and AWS cloud except for GCP it will be performed on an existing mount based indexing.
- Host consistent snapshot is supported for EXT2 file system only if it is mounted as read-only.
- 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.
- To restore files and folders from VM snapshot recovery points, the `/etc/fstab` file on the Linux servers, must have entries based on the file system UUID, instead of device paths. The device paths can change depending on the order in which Linux discovers the devices during system boot.

- While restoring application or file systems from one OS version to another OS version, refer to the OS and application vendor's compatibility matrix. Restore of file system from higher version to lower version is not recommended.
- 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.
- The agentless connection cannot restore the encrypted file by Windows (or EFS) through Granular File Level Restore (Restore Files and Folder option). However you can restore the file through volume level restore and then decrypt the file.

Restoring files and folders from cloud virtual machines

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

Note: For Microsoft Azure, Google Cloud Platform, and Amazon AWS NetBackup supports snapshot and recovery of cloud assets that are encrypted using the keys that 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**.
- 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 or two backup from snapshot job for the assets are complete, 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.

Performing steps after volume restore containing LVM

You can perform steps after volume restore for the LVM volumes.

Note: SFR (Single File Restore) or GRT (Granule Restore) and application restore is performed through the installed agents. But for volume recovery, it is necessary to make the associated file systems online after successful recovery.

To perform steps after volume restore

- 1 Run the command to see all newly attached post volumes on to the host.`pvs`

If there are duplicate PVs (a warning is displayed on the above command) then run,

```
vgimportclone --import /dev/<Device1> /dev/<Device2> ...
--basevgname <NewVGName>
```

Else, find out the newly created Volume Groups (VG) on the host. If new VGs are not displayed then import the VG using the following command. It will discover new VG as <NewVGName>

```
vgimport -a

vgs
```

- 2 Run below command to list all the logical volumes (new and old)

```
lvs <NewVGName>
```

- 3 Activate all the LVs belonging to <NewVGName> as,

```
lvchange --activate y /dev/mapper/<NewVGName>-<LVName1>

lvchange --activate y /dev/mapper/<NewVGName>-<LVName2>

lvchange --activate y /dev/mapper/<NewVGName>-<LVNameN>
```

- 4 Identify the UUID and file system of an authenticate and newly activated LV.

```
blkid -p /dev/mapper/<NewVGName>-<LVName1>
```

```
Output: /dev/mapper/<NewVGName>-<LVName1>:
UUID="2a4bdc14-b5eb-4ee6-b876-ebdcb66c55d9"
BLOCK_SIZE="4096"TYPE="xfs"  USAGE="filesystem"
```

```
blkid -p /dev/mapper/<OldVGName>-<LVName1>
```

```
Output: /dev/mapper/<OldVGName>-<LVName1>:
UUID="2a4bdc14-b5eb-4ee6-b876-ebdcb66c55d9"
BLOCK_SIZE="4096"TYPE="xfs"  USAGE="filesystem"
```

- 5 If the UUID is the same, then you need to change it as follows

File System	Steps
xfs	<pre>mkdir <NewMountPoint> mount -o nouuid /dev/mapper/<NewVGName>-<LVName1> <NewMountPoint> umount <NewMountPoint> xfs_admin -U generate /dev/mapper/<NewVGName>-<LVName1> mount /dev/mapper/<NewVGName>-<LVName1> <NewMountPoint></pre>
ext2 / ext3/ ext4	<pre>mkdir<NewMountPoint> tune2fs -U random /dev/mapper/<NewVGName>-<LVName1> mount /dev/mapper/<NewVGName>-<LVName1> <NewMountPoint></pre>

- 6 If the UUID is different, then run the following command.

```
mount /dev/mapper/<NewVGName>-<LVName1> <NewMountPoint>
```

Troubleshooting

Troubleshooting snapshot restore process for Microsoft Azure 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.

Filtering unsupported files and folders

If you try to restore files or folders from partition or file system that are not supported by CloudPoint, then you get the following error in the restore job.

```
Error nbcs (pid=<processss id>) Failed to restore file(s) and folder(s)
from snapshot for asset <asset name>
```

Workaround:

If you want to avoid listing files or folder while browsing for single file restore, which are not supported by CloudPoint, then enable CP DISKMAP check by setting below flag in `bp.conf` file of NetBackup master server.

`CP_DISKMAP_CHECK = true/yes`

Troubleshooting protection and recovery of cloud assets

This chapter includes the following topics:

- [Troubleshoot cloud workload protection issues](#)
- [Troubleshoot PaaS workload recovery 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 *NetBackup CloudPoint Install and Upgrade Guide*.

Log files for configuration

Use the following logs to troubleshoot cloud configuration issues.

Table 6-1 Log files for configuration

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

Log files for asset discovery

Use the following logs to troubleshoot asset discovery issues.

Table 6-2 Log files for asset discovery

Process	Logs
ncfnbcs Verifies if discovery was completed or not.	Windows <i>NetBackup install path/bin/vxlogview -o 366</i> UNIX <i>/usr/opensv/netbackup/bin/vxlogview -o 366</i>
Picloud Provides the details of discovery operation.	Windows <i>NetBackup install path/bin/vxlogview -i 497</i> UNIX <i>/usr/opensv/netbackup/bin/vxlogview -i 497</i>

Table 6-2 Log files for asset discovery (*continued*)

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

Log files for snapshot creation

Use the following logs to troubleshoot snapshot creation issues.

Table 6-3 Log files for snapshot creation

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

Table 6-3 Log files for snapshot creation (*continued*)

Process	Logs
nbrb	Windows
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.	<p><i>NetBackup install path/bin/vxlogview -o 118</i></p> <p>UNIX</p> <p><i>/usr/opensv/netbackup/bin/vxlogview -i 118</i></p>

Log files for restore operations

Use the following logs to troubleshoot restore issues.

Table 6-4

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

Log files for snapshot deletion

Use the following logs to troubleshoot snapshot deletion issues.

Table 6-5 Log files for snapshot deletion

Process	Logs
bpdm The snapshot delete or clean-up operation is triggered by bpdm.	Windows <i>NetBackup install path/netbackup/logs</i> UNIX <i>/usr/opensv/netbackup/logs/bpdm</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>

Pre-recovery check fails with access denied error during alternate location restore

When attempting to perform recovery of a VM from a backup image copy, if you do not have the required privileges assigned to your role to perform alternate location restore, you encounter the error during the pre-recovery check operation.

This may happen when you have privilege to perform only original location recovery, and you are trying to do alternate location recovery.

Workaround

- While doing original location restore, do not change any pre-populated fields in the pre-recovery page.
- If you want to perform alternate location recovery, ensure that you have the required privileges.

Troubleshoot PaaS workload recovery issues

Error 150: Termination requested by administrator

Explanation: This appears when you manually cancel a restore job from the activity monitor and a database is created on the portal during the partial restore operation.

Workaround: Manually cleanup the database on the provider portal.

Stale status messages in Activity monitor

Explanation: If the CloudPoint container service restarts abruptly; the provider protected restore jobs may remain in the active state and you may not see the updated status on the activity monitor details page.

Workaround: Restart the workflow containers using the following command in the CloudPoint server:

```
docker restart flexsnap-workflow-system-0-min
flexsnap-workflow-general-0-min
```

After restarting the containers, the restore jobs are updated with the latest status in the activity monitor.