

NetBackup™ Web UI Cloud Administrator's Guide

Release 10.4

Last updated: 2024-03-27

Legal Notice

Copyright © 2024 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	Managing and protecting cloud assets	7
	About protecting cloud assets	8
	Limitations and considerations	10
	Configure Snapshot Manager in NetBackup	11
	Add a Snapshot Manager	12
	Add a cloud provider for a Snapshot Manager	12
	Associate media servers with a Snapshot Manager	18
	Discover assets on Snapshot Manager	19
	Enable or disable a Snapshot Manager	20
	(Optional) Add the Snapshot Manager extension	20
	Managing intelligent groups for cloud assets	21
	Considerations for cloud intelligent groups	21
	Create an intelligent group for cloud assets	22
	Delete an intelligent group for cloud assets	26
	Protecting cloud assets or intelligent groups for cloud assets	26
	Customize or edit protection for cloud assets or intelligent groups	29
	Remove protection from cloud assets or intelligent groups	29
	Cloud asset cleanup	29
	Cloud asset filtering	30
	AWS and Azure government cloud support	34
	About protecting Microsoft Azure resources using resource groups	34
	Before you begin	35
	Limitations and considerations	35
	About resource group configurations and outcome	36
	Troubleshoot resource group permissions	39
	About the NetBackup Accelerator for cloud workloads	39
	How the NetBackup Accelerator works with virtual machines	40
	Accelerator forced rescan for virtual machines (schedule attribute)	41
	Accelerator backups and the NetBackup catalog	41
	Accelerator messages in the backup job details log	41
	Configuring backup schedules for cloud workloads	42
	Backup options for cloud workloads	45
	Snapshot replication	49

Configure AWS snapshot replication	49
Using AWS snapshot replication	52
Support matrix for account replication	54
Protect applications in-cloud with application-consistent snapshots	56
Protecting AWS or Azure VMs for recovering to VMware	57
Protecting PaaS assets	58
Prerequisites for protecting PaaS assets	58
Installing the native client utilities	61
Configuring the storage server for instant access	67
Prerequisites for protecting Amazon RDS SQL Server database assets	68
Configuring storage for different deployments	69
About incremental backup for PaaS workloads	71
About archive redo log backup for PaaS workloads	71
Limitations and considerations	72
Discovering PaaS assets	80
Viewing PaaS assets	81
Managing PaaS credentials	81
View the credential name that is applied to a database	82
Add credentials to a database	82
Add protection to PaaS assets	88
Perform backup now	89

Chapter 2 Recovering cloud assets 91

Recovering cloud assets	91
Perform rollback recovery of cloud assets	98
Recovering AWS or Azure VMs to VMware	98
Post-recovery considerations for cloud VMs recovered to VMware	100
Steps to recover images from cloud VMs to VMware	100
Recovering PaaS assets	103
Recovering non-RDS PaaS assets	104
Recovering RDS-based PaaS asset	105
Recovering Azure-protected assets	106
Recovering duplicate images from AdvancedDisk	108

Chapter 3 Performing granular restore 110

About granular restore	110
Supported environment list	111
List of supported file systems	112
Before you begin	113

Limitations and considerations	115
Restoring files and folders from cloud virtual machines	118
Restoring volumes on cloud virtual machines	121
Performing steps after volume restore containing LVM	122
Troubleshooting	124

Chapter 4

Troubleshooting protection and recovery of cloud assets	130
Troubleshoot cloud workload protection issues	130
Error Code 9855: Error occurred while exporting snapshot for the asset: <asset_name>	135
Backup from snapshot jobs take longer time than expected	135
Backup from snapshot job fails due to connectivity issues when Snapshot Manager is deployed on an Ubuntu host	135
Error disambiguation in NetBackup UI	136
Troubleshoot PaaS workload protection and recovery issues	136
Troubleshooting Amazon Redshift issues	142

Managing and protecting cloud assets

This chapter includes the following topics:

- [About protecting cloud assets](#)
- [Limitations and considerations](#)
- [Configure Snapshot Manager in NetBackup](#)
- [Managing intelligent groups for cloud assets](#)
- [Protecting cloud assets or intelligent groups for cloud assets](#)
- [Cloud asset cleanup](#)
- [Cloud asset filtering](#)
- [AWS and Azure government cloud support](#)
- [About protecting Microsoft Azure resources using resource groups](#)
- [About the NetBackup Accelerator for cloud workloads](#)
- [Configuring backup schedules 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](#)
- [Protecting AWS or Azure VMs for recovering to VMware](#)
- [Protecting PaaS assets](#)

About protecting cloud assets

Using NetBackup, you can now protect your in-cloud workloads. The cloud data protection framework leverages the Snapshot Manager infrastructure to drive faster proliferation of cloud providers. In NetBackup 8.3 and later, Snapshot Manager can protect assets in AWS, Azure, Azure Stack hub, GCP, and from version 10.4 OCI cloud is also supported.

The following table describes the tasks.

Table 1-1 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 the Manage Assets and Manage Protection Plans permissions.</p>
Deploy Snapshot Manager	<p>Install Snapshot Manager in your environment.</p> <p>See "Add a Snapshot Manager" on page 12.</p> <p>Review Snapshot Manager and NetBackup limitations.</p> <p>See "Limitations and considerations" on page 10.</p>
	<p>Register the Snapshot Manager in NetBackup.</p> <p>See the <i>NetBackup Snapshot Manager Install and Upgrade guide</i>.</p>

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

Task	Description
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 Snapshot Manager” on page 12.</p>
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 the case of applications, you can set a discovery interval between 15 minutes to 45 minutes.</p> <p>See “Discover assets on Snapshot Manager” on page 19.</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 49.</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 the application-consistent snapshots. See “Protect applications in-cloud with application-consistent snapshots” on page 56.</p>

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

Task	Description
Recover cloud assets	<ul style="list-style-type: none">You can recover the assets using the recovery points. See “Recovering cloud assets” on page 91. See “Recovering cloud assets” on page 91. See “Perform rollback recovery of cloud assets” on page 98.You can also restore the assets using the <code>nbcloudrestore</code> CLI utility. Note: Do not use the <code>bprestore</code> CLI for restores. See the NetBackup Commands Reference Guide.
Troubleshooting	See “Troubleshoot cloud workload protection issues” on page 130.

Limitations and considerations

Consider the following when protecting cloud workloads

- Deletion of Snapshot Manager 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 Snapshot Manager images that are associated with that plug-in.
- Review the *NetBackup Snapshot Manager Install and Upgrade Guide* for information on the capabilities of Snapshot Manager.
- If you have a previous installation of Snapshot Manager, Veritas recommends that you upgrade the Snapshot Manager and not reinstall it.
If you do reinstall the Snapshot Manager server, you need to reconfigure the Snapshot Manager and perform all the protection-related steps.
- By default, Snapshot Manager is configured with port 443.
- After the Snapshot Manager 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 Snapshot Manager, enhanced auditing is not supported. Thus, when you add or update a Snapshot Manager, with non-root but NetBackup Admin rights, during auditing the user is shown as root.

- If you deploy Snapshot Manager using the CloudFormation template, when you register the on-host agent with the Snapshot Manager node using the command, the IP address used must be private IP and not public IP.

Note: Veritas recommends having swap space enabled on NetBackup primary servers that would be used to run backup from snapshot jobs for cloud asset groups. The recommended size for swap space must be greater than or equal to 1.5 times the system memory. In scenarios where swap space enablement is not available, it is recommended to have systems with higher memory configuration.

Configure Snapshot Manager in NetBackup

You can add a Snapshot Manager using the NetBackup Web UI. Starting with 8.3, the Snapshot Manager can discover cloud assets on Amazon Web Services, Azure, Azure Stack, GCP and from version 10.4 OCI cloud is also supported.

Consider the following important points:

- You can associate multiple Snapshot Managers to a NetBackup primary server. But you can associate only one Snapshot Manager to one NetBackup primary server.
- You can now manage Snapshot Manager and control the discovery of assets from the NetBackup WebUI, REST API, and CLI without interacting with the Snapshot Manager interfaces.
- For backup from snapshot jobs, the NetBackup media storage associated servers are used instead of Snapshot Manager associated media servers. The NetBackup media storage associated servers must be connected to the Snapshot Manager to facilitate all the Snapshot Manager related operations.

The following table describes the underlying tasks.

Table 1-2 Configuring Snapshot Manager

Task	Description
Add a Snapshot Manager	See “Add a Snapshot Manager” on page 12.
Add cloud providers	To discover assets on the Snapshot Manager, you must add the cloud providers. See “Add a cloud provider for a Snapshot Manager” on page 12.

Table 1-2 Configuring Snapshot Manager (*continued*)

Task	Description
Discover assets on Snapshot Manager	You can discover assets on the Snapshot Manager. See “Discover assets on Snapshot Manager” on page 19.
Associate media servers	To offload snapshots and restore workflows to a media server, you must associate the media server with the Snapshot Manager. See “Associate media servers with a Snapshot Manager” on page 18.

Add a Snapshot Manager

You can add a Snapshot Manager using NetBackup WebUI.

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

To add a Snapshot Manager

- 1 On the left, click **Workloads > Cloud**.
- 2 Click on the **Snapshot Managers** tab.
- 3 Click **Add**.
- 4 In the **Snapshot Manager** field, enter one of the following:
 - The host name or IP address of the Snapshot Manager.
The host name or IP address must be the same as the one you provided at the time of Snapshot Manager configuration during Snapshot Manager installation.
 - If the DNS server is configured, enter the FQDN of the Snapshot Manager.
- 5 In the **Port** field, enter the port number for the Snapshot Manager.
The default port value is 443.
- 6 Click **Save**.

Add a cloud provider for a Snapshot Manager

You can protect the assets on the Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, Microsoft Azure Stack Hub, and Oracle Cloud

Infrastructure (OCI) providers. Starting with 9.0, the Snapshot Manager can discover Amazon Web Services and Microsoft Azure US Government cloud workloads.

To add a cloud provider for Snapshot Manager

- 1** On the left, click **Workloads > 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 **Snapshot Manager**.

5 Enter the required details.

Cloud provider	Parameter	Description
Microsoft Azure	Credential type: Application service principal	
	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.
	Credential type: System managed identity	Enable system-managed identity on Snapshot Manager host in Azure. Note: Assign a role to the system-managed identity.
	Credential type: User managed identity	
	Client ID	The ID of the user-managed identity connected to the Snapshot Manager host.
	<i>The following parameters are applicable for all the above credential types</i>	
	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 with any resource groups.

Cloud provider	Parameter	Description
Microsoft Azure Stack Hub	<i>Using AAD:</i>	The endpoint URL in the following format allows Snapshot Manager to connect with your Azure resources. <code>https://management.<location>.<FQDN></code>
	Azure Stack Hub Resource Manager endpoint URL	
	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>	The endpoint URL in the following format that allows Snapshot Manager to connect with your Azure resources. <code>https://management.<location>.<FQDN></code>
	Azure Stack Hub Resource Manager endpoint URL	
	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.

Cloud provider	Parameter	Description
Amazon AWS	Access Key	The access key ID, when specified with the secret access key, authorizes Snapshot Manager to interact with the AWS APIs.
	Secret Key	The secret key of the application.
	Note: If the Snapshot Manager is configured with IAM Config, the Access Key and Secret Key options are not available.	
	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.
Google Cloud Platform	VPC Endpoint	First DNS name of AWS Security Token Service (STS) endpoint service with no zone specified.
	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.
	Regions	A list of regions in which the provider operates.

Cloud provider	Parameter	Description
Oracle Cloud Infrastructure	Credential type: API Key	
	User OCID	User's OCID for which you generate the credentials.
	Tenancy	Tenant ID of the OCI account.
	Fingerprint	The fingerprint that you obtain while generating the credential.
	Private Key	The private key that you obtain while generating the credential.
	Regions	One or more OCI regions in which you want to discover the cloud assets.
	Credential type: IAM	NetBackup Snapshot Manager must be a part of a dynamic group and that dynamic group must have enough permissions.

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 Snapshot Manager is deployed in cloud, AWS configuration can be configured to use the IAM role for authentication.

See [“Add a cloud provider for a Snapshot Manager”](#) on page 12.

Before proceeding, ensure that the IAM role is configured within AWS. See the *NetBackup Snapshot Manager Install and Upgrade Guide* for details.

The following implementations of the IAM role are supported:

- **Source account:** In this case, the cloud assets that need to be protected are in the same AWS account as Snapshot Manager. 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 Snapshot Manager. Thus, you need to enter the target account and the target role name details along with the region so that Snapshot Manager 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.

IAM Role for OCI Configuration

If the Snapshot Manager is deployed in cloud, OCI configuration can be configured to use the IAM role for authentication.

See [“Add a cloud provider for a Snapshot Manager”](#) on page 12.

Before proceeding, ensure that the IAM role is configured within OCI. See the *NetBackup Snapshot Manager Install and Upgrade Guide* for details.

OCI supports the Source account implementation of the IAM role. Here, the cloud assets that need to be protected are in the same OCI account as Snapshot Manager. Thus, OCI cloud is aware of the OCI account ID and role name, you need to only select the region.

Associate media servers with a Snapshot Manager

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

To associate media servers with a Snapshot Manager

- 1 On the left, click **Workloads > Cloud**.
- 2 Click on the **Snapshot Managers** tab.
- 3 From the menu next to the Snapshot Manager, click **Advanced settings**.
- 4 In the **Media server** tab, select one or more media servers that you want to associate with the Snapshot Manager.
- 5 Click **Save**.

Discover assets on Snapshot Manager

After you configure your cloud providers with a Snapshot Manager, automatic discovery is triggered to discover assets from the cloud. During periodic discovery, NetBackup pulls the asset data from Snapshot Manager every two hours whereas Snapshot Manager pulls the asset data from cloud provider configurations every one hour. If you disable a Snapshot Manager, 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 Snapshot Manager to fetch the assets data available on the Snapshot Manager.

After the first full discovery, NetBackup subsequently performs periodic incremental discovery of assets for the configured Snapshot Manager. It only detects the changes, such as the 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 primary server and the Snapshot Manager, 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 Snapshot Manager level, which does not discover the assets from the Cloud, but only fetches the point-in-time data from Snapshot Manager.

To discover assets on Snapshot Manager

- 1 On the left, click **Workloads > Cloud**.
- 2 Click on the **Snapshot Managers** tab.
- 3 From the menu next to the Snapshot Manager, 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 **Workloads > Cloud**.
- 2 Click on the **Snapshot Managers** tab.
- 3 Click the Snapshot Manager 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 Snapshot Manager assets.

Change the autodiscovery frequency for Snapshot Manager

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.

Enable or disable a Snapshot Manager

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

To enable or disable a Snapshot Manager

- 1 On the left, click **Workloads > Cloud**.
- 2 Click on the **Snapshot Managers** tab.
- 3 Based on the Snapshot Manager status, select **Enable** or **Disable**.

Note: After disabling a Snapshot Manager 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 Snapshot Manager extension

The Snapshot Manager extension serves the purpose of scaling the capacity of the Snapshot Manager host to service a large number of requests concurrently running on the Snapshot Manager server at its peak performance capacity. You can install one or more Snapshot Manager 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 Snapshot Manager host.

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

Supported Snapshot Manager extension environments:

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

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

Managing intelligent groups for cloud assets

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, and PaaS assets 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 a 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 the 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 to 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*.

Considerations for cloud intelligent groups

Consider the following before creating cloud intelligent groups:

- The values that you specify for the intelligent group filters are case-sensitive.
- The **Status** attribute is derived from **State**. To add a condition filter on the **Status** attribute, select **State** from the **Filter** drop-down.
- Account ID option in intelligent groups:
 - The **All Accounts** option in the **Account ID** list is available to NetBackup's Default Cloud Administrator role.
 - The **All Accounts** option in the **Account ID** list is available to NetBackup's custom roles with the **All cloud assets** permission for one or multiple cloud service providers.

- Any custom role with explicit asset access permission of account(s) or subscription(s) cannot use the **All Accounts** option.

Considerations for PaaS intelligent groups

- You can subscribe the assets to different protection plans based on the backup types that are supported for the assets. However, you cannot subscribe the intelligent groups containing AWS RDS Oracle assets to the protection plans that contain incremental schedules.
- The Service type drop-down shows the service types available for the provider, irrespective of the discovered assets.
- Intelligent groups for PaaS assets support protection of the Azure, AWS, and GCP assets.
- Tag handling for Azure SQL server and Azure Managed Instance:
 - For SQL server, the "Server" keyword is added as a prefix to tags, when the tag is copied at the database level.
 - For Azure Managed Instance, the "Instance" keyword is added as a prefix to the tags, when the tag is copied at the database level.
 - No prefix is added for tags in other workloads.

Create an intelligent group for cloud assets

To create an intelligent group for cloud assets

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

Note: If the region is not specified, then the cloud intelligent group filter applies to the assets from all the discovered regions.

- 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 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 groups for cloud assets”](#) on page 24.

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 a 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 has a header with 'AND' and 'OR' buttons. The query is built using three conditions connected by 'OR' logic:

- Condition 1: displayName (keyword) Contains (operator) CP (value)
- Condition 2: tagName (keyword) Starts with (operator) eng (value)
- Condition 3: state (keyword) = (operator) running (value)

Each condition row has a trash icon to remove it. At the bottom right, there are buttons for 'Cancel', 'Add and Protect', and 'Add'.

This example uses **OR** 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: The special character '`<`' is not supported in a tag name. If present, asset group creation fails.

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 by 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 groups for cloud assets

Note: The attribute values may not match exactly with the 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 1-3 Query keywords

Keyword	Description
	(All values are case-sensitive)
<code>displayName</code>	Asset's display name.
<code>state</code>	For example, running, stopped etc.

Table 1-3 Query keywords (*continued*)

Keyword	Description (All values are case-sensitive)
tag	A label assigned to the asset for categorization.
instanceType / machineType / vmSize	Asset's instance/machine type or VM size, depending on the cloud provider selection. For example, t2.large, t3.large, or b2ms, d2sv3
parentEntityName	Name of the parent entity of the asset.
parentEntityType	Entity type of the parent entity of the asset.
resourceGroup	Resource group of the asset.
entityType	Entity type of the asset.
shape	Size of the VM in OCI. For example, VM.Standard.E4.Flex, VM.Standard.E5.Flex, and so on.
compartmentId	The asset's compartmentId. OCI uses the compartmentId to organize and isolate the cloud resources.

Table 1-4 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 group for cloud assets

To delete an intelligent group for cloud assets

- 1 On the left, click **Workloads > Cloud**.
- 2 Locate the intelligent group in 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 groups for cloud assets

You can create cloud provider-specific protection plans for your cloud workloads. Then you can subscribe the assets that are associated with the 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 is automatically converted to the new provider-specific format. This conversion happens after an upgrade to NetBackup 9.1. For example, if you had the assets from Google Cloud and AWS Cloud that are subscribed to one protection plan, then the protection plan is split. The protection plan is split into two separate protection plans for each provider.

See [the section called “Conversion of protection plans after an upgrade to NetBackup 9.1 and later”](#) on page 27. 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 **Workloads > 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**
 For more information about storage options in the web UI, review the *Configuring storage* section in the [NetBackup Web UI Administrator's Guide](#).
 - **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 an upgrade to NetBackup 9.1 and later

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

- Protection plan conversion starts when the asset migration is completed after the upgrade of NetBackup to 9.1 and later.
- Old protection plans with no assets subscribed are not 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.
- The **Backup now** operation fails 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. The protection plan name is updated to retain the asset subscription for any one cloud provider upon conversion.
For the other cloud providers that 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 continues with the assets that are successfully unsubscribed. The conversion process for the assets that failed, is retried every 4 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 when the required access permissions are assigned 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 the 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 **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 remove the protection for.
- 3 Click **Remove protection > Yes**.

Cloud asset cleanup

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

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

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

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

For example,

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

User can also manually run the POST query using the `cleanup-assets` named query with the following request body and then run GET with query ID obtained from the 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'"
    }
  }
}
```

Cloud asset filtering

You can define custom filters based on attributes, which NetBackup can use to list assets in the Virtual machines, Applications, PaaS, and Volumes tabs.

To create a filter

- 1
- On the left, click **Workloads > Cloud**.
- 2
- Under the Virtual machines, Applications, PaaS, or Volumes tab, click on the **Filter** icon on the right top of the screen.
The **Create filter** option is displayed.
- 3
- Click the **Create filter** option to define a custom filter based on attributes to list assets. in the Virtual machines, Applications, PaaS, or Volumes tab.
- 4
- To create a filter, enter the details for the following parameters:

Parameter	Description
Name	Name of the filter.
Description	Describe the filter.
Query	To select only the assets that meet specific conditions, create your query.

- 5 To select only the assets that meet specific conditions, create your own query:
Click **+ condition**.
- 6 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 groups for cloud assets”](#) on page 24.

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:

Create filter

Name *

aws-cloud-assets

Description

Enter description

Query

AND OR

+ Condition + Sub-query

Provider

Contains

aws

Name

Contains

cloudpoint

Cancel

Save

Save and add another

Save and apply

This example uses **AND** to narrow the scope of the query: It selects only the assets that have `aws` in their display name and that also have a **Name** as `cloudpoint`, **and are in** running state.

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.

Query options for creating filter

Note: The attribute values may not match exactly with the 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 1-5 Query keywords

Keyword	Description
	(all values are case-sensitive)
Server type	Type of the server.

Table 1-5 Query keywords (*continued*)

Keyword	Description
	(all values are case-sensitive)
Instance ID	Asset's instance ID, depending on the cloud provider selection.
Instance name	Asset's instance name, depending on the cloud provider selection.
Name	Asset's display name.
Provider	Asset's cloud provider name.
Region	Asset's cloud provider region name.
Config ID	Asset's config ID.
Database service	Asset's database service.
Deleted	Deleted asset.
Entity type	Asset's entity type.
Service domain	Asset's service domain.
Snapshot Manager	The instance of Snapshot Manager with which the asset is registered.

Table 1-6 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.

AWS and Azure government cloud support

Starting with 8.3, the Snapshot Manager can discover Amazon Web Services and Microsoft Azure US Government cloud workloads. After the Snapshot Manager is added to NetBackup, you can protect the workloads by NetBackup. NetBackup is compliant with the regulatory requirements including IPv6 support to deploy Snapshot Manager 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 the provided region. The discovered assets are displayed in NetBackup. Currently, only workloads from selected regions and mapped endpoints are discovered and protected. For the same Snapshot Manager host, you cannot use a combination of public and government clouds.

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

Snapshot Manager 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

Note: PaaS assets does not support government cloud.

For information about configuring AWS and Microsoft Azure, See [“Add a cloud provider for a Snapshot Manager”](#) on page 12.

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 with a resource group. After a snapshot is created, it is associated with a resource group. Also, each resource group is associated with a region. See the following:

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

Snapshot Manager creates a snapshot and places the snapshot in the 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 a 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 snapshot creation succeeds or fails.

To enable this feature, you must create peer resource groups. Snapshot Manager 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, hyphens, or parentheses 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 1-7 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 the resource group of the resource.
Specified	Not 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 the conditions are not met, snapshot jobs fail.</p>
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 with 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 1-8 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 the resource group of resource. 	<ul style="list-style-type: none"> ■ Resource Group Prefix value is provided. ■ The Protect assets even if prefixed Resource Groups are not found check box is selected. 	Snapshots are created in the peer resource group.
<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 the 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 the original resource group not the peer resource group.

Table 1-8 Example configurations (*continued*)

Conditions	Configurations	Result
<ul style="list-style-type: none"> ■ OS and all disks are in separate resource groups, RG1 and RG2. ■ Peer resource group 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 the 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.
<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 with resource groups.

Workaround:

To resolve this issue, perform the following steps:

1. Navigate to <https://portal.azure.com/#blade/HubsExtension/BrowseResourceGroups>.
2. Click on the resource group, that is to be used in the snapshot.
3. Click on **Access control (IAM)**.
4. Click on **Add Role Assignment**.
5. Select **Role as Owner**, **Assign Access to as User**, and select the **Application (created for Snapshot Manager, to make API calls)**.
6. Save and try to back up 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 full backups faster than traditional backups. 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 Snapshot Manager.
- Reduces the CPU load on the Snapshot Manager.

How the NetBackup Accelerator works with virtual machines

For Azure and Azure Stack backups, Accelerator is activated when you select an 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 the **Keep backup only** or **Initiate backup when snapshot is about to expire** option in the protection plan, the snapshot is retained for accelerator purposes 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 such a retention period for the snapshot, it leads to more than 200 incremental snapshots. Then, full backups take place instead of the 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 schedules for cloud workloads

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

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

To add a 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, and 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 (For Azure SQL server, GCP SQL Server, and SQL Managed Instance PaaS assets only.) If you have selected **Protect PaaS assets only** for the protection plan, select **Backup type** as **Incremental backup** or **Full**. For the incremental backup type, NetBackup takes an initial full backup, and all subsequent backups capture only incremental changes in the database. This feature increases backup performance to a great extent. In case of a schema change, goes back to a full backup from an incremental backup, and notifies this activity in the activity monitor.

Assign a longer retention period to full backups than to incremental backups within a policy. A complete restore requires the previous full backup plus all subsequent incremental backups. It may not be possible to restore all the files if the full backup expires before the incremental backups. See [“About incremental backup for PaaS workloads”](#) on page 71.

- 5 From the **Recurrence** drop-down, specify the frequency of the backup.
- 6 In the Snapshot and backup options, do any of the following:
 - Select the **Keep snapshot along with backup** option to retain both the snapshot and the backup. Specify a 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.
 (For Amazon RDS Oracle assets only) If you have selected the **Protect PaaS assets only** option for the protection plan, you can select Backup type as **Full**, **Differential incremental**, or **Archived redo log**.
 For incremental and Archived redo log backup types, NetBackup takes an initial full backup and captures all the subsequent backups by incremental or archive changes in the database. This feature increases backup performance to a great extent.
 Do not use multiple protection plans having incremental schedules. Also, do not use protection plans with a frequency greater than 24 hours in the archived log schedule. For a successful restore NetBackup requires the previous full backup, all subsequent incremental backups, and all subsequent archive backups. If the full backup expires before the incremental or archive backups, it may not be possible to restore all the files.
 - Select **Keep snapshot only** option, to retain only the snapshot. Specify the retention period for the snapshot using the **Keep snapshot for** drop-down.
 - (Optional) If you have selected the 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 49.

- 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 the **Keep backup only** option, to retain only the backup. The snapshot expires immediately after the backup. Specify a retention period for the backup using the **Keep backup for** drop-down. Select **Full** from the **Backup type** drop-down.
- 7 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 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 Snapshot Manager. VxMS-based indexing is currently supported for GCP, 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 and backup 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 Snapshot Manager.
- To index files and folders with the same mount path as mentioned in `/etc/fstab`, the `/etc/fstab` file on the Linux servers must have entries based on the UUID file system, instead of device paths. The device paths can change depending on the order in which Linux discovers the devices during system boot.

Note: If the VM is not in a 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

Note: For a connected VM, a file system consistent snapshot is attempted. In case a connected VM is stopped later, then the application enters into an error state and a crash-consistent snapshot is taken instead of a file system consistent snapshot. You can refer to the Job monitor and see the logs if the snapshot taken was crash-consistent or file system consistent snapshot.

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 — **Backup options** — Permissions — 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.

Create protection plan

Basic properties Schedules Storage options **Backup options** Permissions Review

Backup options

☐ Enable granular recovery for files or folders
☒ Specify snapshot destination resource group

Configuration name *
 azurecloudplugin

Fetching subscription and resource group details may take some time depending upon the network connectivity.

Subscription name or ID *
 XXXXXX (a332d749-XXXXXX-XXXXXX-XXXXXX)

Resource group
 azure-scale-rhel83-mongo-dnd

Region
 eastus2

Select

Cancel Previous Next

Excluding selected disks from backups

You can configure a protection plan to exclude some disks from the backup and snapshot which is 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.

Create protection plan

Basic properties Schedules **Backup options** Permissions Review

Backup options

☐ Enable granular recovery for files or folders
☐ Enable regional snapshot
☒ Exclude selected disks from backups

☐ All non boot disks
☒ Disks with a specific tag

Key
 example_tag_exclude

Value
 example: true

Note: A protection plan that has the 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:

- For LVMs, if disks are excluded partially then the system might not boot up properly.
- If there is a non-supported file system configured on a disk and user wants to exclude that disk from a snapshot. The snapshot continues to be a crash-consistent snapshot as the disk containing the non-supported file system is excluded.
- To exclude the disk, have the **nofail** flag attached to the data disk before taking a snapshot in the `/etc/fstab` file. This is required if the you 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**
- Ensure that the assets are properly discovered after making any changes 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 you attach a tag while the snapshot is in progress, NetBackup does not consider that **tag** as a part of exclusion. Once discovery is complete, it will be considered for the next protection cycle.
- In case of an OS with a non-English locale, if you opt for tag based exclusion in the protection plan and the disk tags contain non-English characters, even then, disk exclusion works as expected. But in some cases, a tag with non-English characters is not correctly captured in the 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 regions within the same account.
- same region in a different accounts.
- different regions within different accounts.

For example, AWS cloud administrators 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 the same account or X/Y region in a different account, for an added level of protection. In NBU Snapshot Manager 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 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 Snapshot Manager. 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 Snapshot Manager.
Additionally, to replicate encrypted snapshots to a cross-account, the encryption CMK key from the original location needs to be shared with 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 the 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 Snapshot Manager's 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 inside the same region, but a different account 'aws_account_2-us-east-1'.
- Assets that are discovered by a Snapshot Manager must be replicated to the region that is discovered by the same Snapshot Manager.
For example, assets that are discovered by Snapshot Manager 'CP1' cannot be subscribed to a protection plan replicating to a region that is discovered by Snapshot Manager '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 of 5, EBS has a limit of 5, and EC2 has a limit of 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 describes how to create snapshot replicas using the AWS snapshot replication feature, and restore the replicated snapshots whenever required. Refer to the *NetBackup Snapshot Manager 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 Snapshot Manager (CP1) in web UI.
See [“Add a Snapshot Manager”](#) on page 12.
- 2 Add AWS plug-in for Source and Target region for replication.
- 3 Create a protection plan and select **Region** and **Account**.
See [“Configuring backup schedules for cloud workloads”](#) on page 42.
- 4 Connect and configure an application consistent guest VM using the OnHost agent.
- 5 Start the 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 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 Snapshot Manager in a failed over region.

Restoring in the target region

- 1 Disable server CP1 in the Source region from web UI.
See [“Enable or disable a Snapshot Manager”](#) on page 20.
- 2 Register a new Snapshot Manager (CP2) in the target region, from web UI.
- 3 Add AWS plug-in for only the Target region and account. Let the discovery complete.

- 4 To restore VMs:
 - Sign in to the NetBackup Web UI.
 - On the left, click **Workloads > Cloud**. On the **Virtual machines** tab, click the computer 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 91.
- 5 Add the appropriate security group to the restored VMs to enable remote access.
- 6 Uninstall and reinstall the Snapshot Manager agent from the restored VMs, and then register the Snapshot Manager 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. Start a snapshot-based backup.

Restoring 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 Start a snapshot-based backup and replicate.
- 4 Disable the CP2 server in web UI. See [“Enable or disable a Snapshot Manager”](#) on page 20.
- 5 Enable the CP1 server and start a deep discovery from the AWS provider console.
- 6 Perform full restore of the VMs from the Target region.
- 7 Add the appropriate security group to enable remote access to the restored VMs.
- 8 Uninstall and reinstall the Snapshot Manager agents from the restored VMs. Then register Snapshot Manager 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 1-9 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 1-10 Support matrix for different accounts in the 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 1-11 Support matrix for different accounts in different regions 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.
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 VMs in the 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 restoration, consider the following:

- For alternate location restore of MS SQL workloads, the target host must be discovered but the application status should not be in a 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 a 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 Snapshot Manager 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 to 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**.

Protecting AWS or Azure VMs for recovering to VMware

NetBackup allows you to protect AWS and Azure VMs and restore the protected VMs as on-premises VMware VMs. This section describes the considerations and prerequisites for the same.

- NetBackup supports recovering backup images from AWS EC2 or Azure VMs in the MSDP storage server and MSDP Cloud, without using Glacier or Archive mode.
- Supported operating systems of the source VMs:
 - Windows Server 2022 Series
 - RHEL 9.x
 - SUSE 15SP5: For source VM in AWS provider, must be created from HVM AMI. NetBackup does not support ARM-type VM conversion.
- The network interface in the source VM must use DHCP and be enabled on boot.
- If the source VM platform is Linux, it must be a UUID Linux VM used in `/etc/fstab`. If the source VM platform is Windows, enable `pagefile` on the C drive.

Protecting PaaS assets

The **Applications** tab displays PaaS RDS assets, while the **PaaS** tab displays non-RDS PaaS assets. You can view, protect, and recover PaaS assets from these two tabs.

Prerequisites for protecting PaaS assets

NetBackup lets you discover, protect, and restore PaaS assets across different cloud platforms for a variety of assets. This section details the supported platforms and databases.

Supported cloud providers

NetBackup enables you to protect PaaS assets with the following cloud providers:

- Microsoft Azure
- AWS
- GCP

Supported databases for different providers

The following table lists the supported databases for each cloud provider.

Table 1-12 Supported databases by PaaS

Providers	Supported databases
Microsoft Azure	PostgreSQL, SQL Managed Instance, SQL, MariaDB, Azure Cosmos DB for NoSQL, Azure Cosmos DB for MongoDB, and MySQL. The following components are not supported: Azure SQL - Elastic pool Azure SQL Managed Instance - Azure Arc Azure Cosmos DB for MongoDB vCore Azure PostgreSQL - Hyperscale (Citus) server group and Azure Arc enabled PostgreSQL Hyperscale
AWS	RDS SQL, RDS PostgreSQL, RDS MySQL, RDS MariaDB, RDS Aurora MySQL, RDS Aurora PostgreSQL, Amazon RDS for Oracle, Amazon Redshift, and DynamoDB.
GCP	Cloud SQL for PostgreSQL, Cloud SQL for SQL Server, and Cloud SQL for MySQL

Supported platforms

This section details the supported platforms for primary and media servers.

Table 1-13 Supported platforms for PaaS

NetBackup server	Supported platform
Primary	RHEL, SUSE, and Windows
Media	RHEL
Storage server	Universal share on underlying MSDP block storage or MSDP-Cloud storage STU

Required cloud provider permissions

The credential that you use to add the cloud providers must have the required permissions and privileges assigned, as mentioned in the *NetBackup Snapshot Manager Installation and Upgrade Guide*.

Supported ports

Here are the supported ports for different PaaS databases. Note that the AWS RDS workloads support custom ports along with the default port.

Table 1-14 Supported ports for PaaS

Database PaaS workload	Supported ports
Azure SQL Server	1433
Azure SQL Managed Instance	1433
Azure MySQL	3306
Azure PostgreSQL	5432
Azure MariaDB	3306
GCP PostgreSQL	5432
GCP MySQL	3306
AWS DynamoDB	NA
AWS RDS PostgreSQL	5432
AWS RDS MySQL	3306

Table 1-14 Supported ports for PaaS (*continued*)

Database PaaS workload	Supported ports
AWS MariaDB	3306
AWS RDS AuroraDB Postgres	5432
AWS RDS AuroraDB MySQL	3306
AWS RDS SQL server	1433
Amazon RDS for Oracle	1521
Azure Cosmos DB for NoSQL	443
Azure Cosmos DB for MongoDB	10255
GCP SQL Server port	1433
Amazon Redshift	5439

Enabling binary logging for MySQL databases

- For AWS, see <https://aws.amazon.com/premiumsupport/knowledge-center/rds-mysql-functions/>
- For Azure, set the value of the parameter `log_bin_trust_function_creators` as 1, as described in the link: <https://learn.microsoft.com/en-us/azure/mysql/single-server/how-to-server-parameters>
- For GCP, do the following:
 - Open the instance and click **Edit**.
 - Scroll down to the **Flags** section.
 - To set a flag, click **Add item**, select `log_bin_trust_function_creators` flag from the drop-down menu, and set its value to on.
 - Click **Save** to save your changes. You can confirm your changes under **Flags** in the **Overview** page.

Enabling backup and restore in Kubernetes

Before performing backup and restore operations on AKS and EKS deployments, you need to configure the parameter `MEDIA_SERVER_POD_CIDR` in the `bp.conf`

file in the primary server pod. Specify its value as the subnets on which the media server pod is deployed. You can use comma-separated values. For example:

```
MEDIA_SERVER_POD_CIDR=10.0.0.0/8, 10.0.0.0/16
```

Installing the native client utilities

If you use a build-your-own (BYO) setup, you must install the native client utilities in your NetBackup environment for your PaaS workload to work. If the BYO setup is configured to use an unauthorized user (or service user) account, ensure that the NetBackup service user has the required execute permissions on the native client utilities.

For NetBackup deployments in Azure Kubernetes Services (AKS) or Elastic Kubernetes Services (EKS), the native client utilities are packaged as part of the NetBackup media server, primary server, and data mover container image. Manual installation is not required for them.

Ensure that the network settings like firewall, security group, and DNS configuration are configured appropriately to access databases within the cloud provider.

Starting with NetBackup 10.4, the DBPaaS agent and the native client utilities run under the service user, if the service user is configured.

Note: If any of these packages are already installed in the media server(s), remove the packages to avoid conflict with the newer versions of the packages that you install.

Installing the MySQL client utility

Note: The MySQL client utility recommended version is 8.0.34.

RPM Download location <https://downloads.mysql.com/archives/community/>

To install, run the following commands in the terminal:

- 1 `rpm -ivh mysql-community-common-<version_no>.x86_64.rpm`
- 2 `rpm -ivh mysql-community-client-plugins- <version_no>.x86_64.rpm`
- 3 `rpm -ivh mysql-community-libs- <version_no>.x86_64.rpm`
- 4 `rpm -ivh mysql-community-client- <version_no>.x86_64.rpm`

Note: Avoid using MySQL client utility 8.0.32 version as there is a bug reported by MySQL.

Installing the sqlpackage client utility

Note: The `sqlpackage` client utility recommended version is 19.2 (Build: 162.0.52).

Download locations <https://docs.microsoft.com/en-us/sql/tools/sqlpackage-download?view=sql-server-ver15>

https://packages.microsoft.com/rhel/7/prod/msodbcsql17-17.9.1.1-1.x86_64.rpm

https://packages.microsoft.com/rhel/7/prod/unixODBC-2.3.7-1.rh.x86_64.rpm

To install, run the following commands in the terminal:

```
1 cd ~
2 mkdir sqlpackage
3 unzip ~/Downloads/sqlpackage-linux-<version string>.zip -d
  ~/sqlpackage
4 echo "export PATH=\"\$PATH:$HOME/sqlpackage\"" >> ~/.bashrc
5 chmod a+x ~/sqlpackage/sqlpackage
6 source ~/.bashrc
```

Note: Ensure that the `sqlpackage` is added as a default path variable. If you still get the: Cannot find `sqlpackage` error, restart the NetBackup services on the media server.

```
7 sqlpackage
8 rpm -ivh unixODBC-2.3.7-1.rh.x86_64.rpm
9 rpm -ivh msodbcsql17-17.10.2.1-1.x86_64.rpm
```

RHEL 9 users perform the following additional steps:

- 1 Download the Microsoft.NETCore.App.Runtime.linux-x64 from the link:
<https://www.nuget.org/api/v2/package/Microsoft.NETCore.App.Runtime.linux-x64/6.0.10>
Locate the file: `microsoft.netcore.app.runtime.linux-x64.6.0.10.nupkg`.
- 2 Extract the file using a decompression tool like, 7zip.
- 3 Navigate to:
`microsoft.netcore.app.runtime.linux-x64.6.0.10.nupkg\runtimes\linux-x64\lib\net6.0\`
- 4 Copy the file `System.Security.Cryptography.X509Certificates.dll` from there to the `~/sqlpackage` folder created in step 2 of installing the *sqlpackage* client utility task.

If you are attaching the 10.1 media server as an external media server with 10.1.1 NetBackup setup, perform the following steps on the 10.1 media server.

For a BYO NetBackup setup:

- Run the command:
`mkdir -p <backup and restore ushare export path>`
- Check the `Defaultvers` value of NFS in the `/etc/nfsmount.conf` file.
 - If the `Defaultvers` value is `nfs3`, then mount the backup and restore ushare path with the option `nolock`. For example: `mount <ushare mount path> <ushare export path> -o nolock`
 - If the `Defaultvers` is `nfs4`, mount the backup and restore ushare path without the `nolock` option.

For NetBackup deployed in AKS and EKS environments:

- Run the command:
`mkdir -p <backup and restore ushare export path>`
- Check the `Defaultvers` value of NFS from the `/etc/nfsmount.conf` file.
 - If the `Defaultvers` value is `nfs3`, then mount the backup and restore ushare path with the option `nolock` For example: `mount <ushare mount path> <ushare export path> -o nolock`
 - If the `Defaultvers` value is `nfs4`, then mount the v4 version of backup and restore ushare path without the `nolock` option.

Installing PostgreSQL client utility

PostgreSQL client utility recommended version is 15.3.

Download locations

RHEL 7	https://download.postgresql.org/pub/repos/yum/15/redhat/rhel-7-x86_64/
RHEL 8	https://download.postgresql.org/pub/repos/yum/15/redhat/rhel-8-x86_64/
RHEL 9	https://download.postgresql.org/pub/repos/yum/15/redhat/rhel-9-x86_64/

To install, run the following commands in the terminal:

```
1 rpm -ivh postgresql15-libs-15.3-1PGDG.rhel7.x86_64.rpm
2 rpm -ivh postgresql15-15.3-1PGDG.rhel7.x86_64.rpm
```

Note: `lz4` compression package and `libcuckoo` are required by `postgresql15-15.3-1PGDG.rhel8.x86_64.rpm` on RHEL 8 and 9.

Installing MongoDB client utility

MongoDB client utility recommended version is 100.7.3.

Download locations

RHEL 7	https://fastdl.mongodb.org/tools/db/mongodb-database-tools-rhel70-x86_64-100.7.3.rpm
RHEL 8	https://fastdl.mongodb.org/tools/db/mongodb-database-tools-rhel80-x86_64-100.7.3.rpm
RHEL 9	https://fastdl.mongodb.org/tools/db/mongodb-database-tools-rhel90-x86_64-100.7.3.rpm

To install, run the following commands in the terminal:

```
rpm -ivh mongodb-database-tools-rhel70-x86_64-100.7.3.rpm
```

Installing the Amazon RDS for Oracle client utility

Amazon RDS for Oracle client utility recommended version is 21.11.0.0.0-1.el8.

Download locations

instantclient-basic	https://download.oracle.com/oh_solware/hw/instantclient/2111000/odbc/instantclient-basic-21.11.0.0.0-1.el8.x86_64.rpm
instantclient-odbc	https://download.oracle.com/oh_solware/hw/instantclient/2111000/odbc/instantclient-odbc-21.11.0.0.0-1.el8.x86_64.rpm

To install, run the following commands in the terminal:

```
1 c. yum install unixODBC
2 rpm -ivh oracle-instantclient-basic-21.10.0.0.0-1.el8.x86_64.rpm
3 d. rpm -ivh oracle-instantclient-odbc-21.10.0.0.0-1.el8.x86_64.rpm
```


Installing the EFS utility and configuring permissions

To install the EFS utility

- 1 Visit this page in AWS documentation:
<https://docs.aws.amazon.com/efs/latest/ug/installing-amazon-efs-utils.html>
- 2 Refer to the section: *To build and install amazon-efs-utils as an RPM package for Amazon Linux, Amazon Linux 2, and Linux distributions other than OpenSUSE or SLES.*
- 3 Install `stunnel` version 5.
- 4 Modify region from `/etc/amazon/efs/efs-utils.conf` to your RDS instance region.

Configuring EFS and the restore mount path on AWS

You need to configure the Amazon Elastic File System (EFS) before you can perform any backup.

- To configure EFS, see the following knowledge-base article:
https://www.veritas.com/support/en_US/article.100059038
- To configure the mount path for restore, see the following knowledge base article:
https://www.veritas.com/support/en_US/article.100059039

Configuring AWS permissions for NetBackup

NetBackup required permissions in AWS to perform backup and restore. To configure the permissions, create an AWS IAM role and assign the permissions required by NetBackup to the role. For more information on how to create an IAM role, see this link in AWS documentation:

<https://docs.aws.amazon.com/iam/index.html>

Required permissions:

```
efsdescribemounttarget:
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "elasticfilesystem:DescribeMountTargets",
        "elasticfilesystem:DescribeFileSystems",
      ],
    }
  ]
}
```

```

        "Resource": [
            "arn:aws:elasticfilesystem:*:*:access-point/*",
            "arn:aws:elasticfilesystem:*:*:file-system/*"
        ]
    }
}

rdsdescribeoptiongroup
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "rds:CreateOptionGroup",
                "rds:ModifyDBInstance",
                "rds:ModifyOptionGroup",
                "s3:*"
            ],
            "Resource": "*"
        }
    ]
}

AmazonRDSReadOnlyAccess:(AWS Managed)
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "rds:Describe*",
                "rds:ListTagsForResource",
                "ec2:DescribeAccountAttributes",
                "ec2:DescribeAvailabilityZones",
                "ec2:DescribeInternetGateways",
                "ec2:DescribeSecurityGroups",
                "ec2:DescribeSubnets",
                "ec2:DescribeVpcAttribute",
                "ec2:DescribeVpcs"
            ],

```

```

        "Resource": "*"
    },
    {
        "Effect": "Allow",
        "Action": [
            "cloudwatch:GetMetricStatistics",
            "cloudwatch:ListMetrics",
            "cloudwatch:GetMetricData",
            "logs:DescribeLogStreams",
            "logs:GetLogEvents",
            "devops-guru:GetResourceCollection"
        ],
        "Resource": "*"
    },
    {
        "Action": [
            "devops-guru:SearchInsights",
            "devops-guru:ListAnomaliesForInsight"
        ],
        "Effect": "Allow",
        "Resource": "*",
        "Condition": {
            "ForAllValues:StringEquals": {
                "devops-guru:ServiceNames": [
                    "RDS"
                ]
            },
            "Null": {
                "devops-guru:ServiceNames": "false"
            }
        }
    }
]
}

```

Configuring the storage server for instant access

Here is the required configuration for your storage server to support instance access.

- 1 Ensure that NFS and NGINX are installed.
- 2 The NGINX version must be the same as the one in the corresponding official RHEL version release. Install it from the corresponding RHEL yum source (EPEL).
- 3 Ensure that the `policycoreutils` and `policycoreutils-python` packages are installed from the same RHEL yum source (RHEL server). Run the following commands:
 - `semanage port -a -t http_port_t -p tcp 10087`
 - `setsebool -P httpd_can_network_connect 1`
- 4 Ensure that any mount points do not directly mount the `/mnt` folder on the storage server. Mount the mount points to its subfolders only.
- 5 Enable the `logrotate` permission in `selinux` using the following command:

```
semanage permissive -a logrotate_t
```

Prerequisites for protecting Amazon RDS SQL Server database assets

You need to enable the native backup and restore options for the option group to protect RDS SQL assets.

This option group must be a part of the IAM role, which has an *AWSBackupServiceRolePolicyForRestores/Backup* policy attached to it.

To create an option group:

- 1 On your AWS portal, go to IAM and create a new role.
- 2 Attach the following permissions:
 - `AWSBackupServiceRolePolicyForRestores`
 - `AWSBackupServiceRolePolicyForBackup`
 - `sqlNativeBackup`
- 3 Go to RDS > Option groups. Do the following:

Note: Refer to the AWS documentation for the latest steps on adding the native backup and restore option to the option group for RDS SQL Server.

- Create a group (Name: `SqlServerBackupRestore`, Description: `xxx`, Engine: Select your database engine, Major Engine Version: Select the version of your DB instance).

- Click Create.
 - Click the created group to edit. Do the following:
 - Click Add option.
 - Select the option: `SQLSERVER_BACKUP_RESTORE`.
 - Select the IAM role you created in the previous steps.
 - Select Immediately, to schedule an instant change.
- 4 Go to RDS > Databases and select your instance. Do the following:
- Click Modify.
 - Select the option group you created in the previous step.
 - Click Next.
 - Select Apply immediately, to avoid service downtime.
 - Click Modify DB instance to apply the changes.
 - If you have a connection to database from the SQL Management Studio, close and connect again.

To use the bucket for protecting Amazon RDS SQL Server database assets:

For NetBackup versions before 10.2:

- 1 Create the NetBackup AWS S3 bucket as:
`netbackup-<AWS_ACCOUNT_IDENTIFIER>`
- 2 If a bucket already exists with the same convention, NetBackup uses it.

For NetBackup versions 10.2 onwards:

- ◆ The bucket is automatically created if it is not present, following the convention:`netbackup-<AWS_ACCOUNT_IDENTIFIER>-region`

Configuring storage for different deployments

This section describes how to configure storage for different NetBackup deployments.

For MSDP cloud deployments

MSDP storage targets use the media servers. Install the native client utility on the media server, and connect that media server to the PaaS workload.

For MSDP cloud volume storage, NetBackup protects PaaS assets through the Data Mover Container (DMC), using the universal share accelerator.

The Universal share accelerator requires a minimum of 500GB storage space, as persistent storage, to store the temporary metadata within the DMC. The storage path must be the same as the one in the MSDP storage server.

For Kubernetes deployments

Consider the following:

- Create the persistent volume claims using disk-based and delete-policy storage classes, and attach them to the container at the storage path.
- It is recommended to use the default storage class with the default storage size of 600 Gi storage. To change the storage class or storage size, you need to update the `pdconf` config map of the Kubernetes deployment, as follows:

```
STORAGE_CLASS=<disk-based storage class>  
STORAGE_SIZE=<pv size>
```

For VM-based BYO deployments

Consider the following:

- Mount a new disk with 600 GB storage in NetBackup Snapshot Manager, at the path `:/datamover_storage`.
- Each datamover container creates a directory at the mounted disk path, and a symlink as the storage path. You can see this path in the datamover container as the storage path. This path is the same as the MSDP storage path for the temporary storage for universal share accelerator operation.

If you do not have sufficient storage space available on the deployment, you can override the storage requirement. Do the following:

1. Navigate to:
`/cloudpoint/openv/netbackup/vpfs_override_parameters.json`
2. Update the `CloudCacheSize` parameter with the available storage size in GBs.

```
{  
  "DataTransferManagementOptions": {  
    "CloudCacheSize": 200  
  }  
}
```

About incremental backup for PaaS workloads

NetBackup supports differential incremental backup for Azure SQL Server, Azure SQL Managed Instance, AWS RDS Oracle, and GCP SQL Server workloads. Incremental backups reduce the backup window significantly in NetBackup. In this method, NetBackup backs up only the data that has been changed since the subsequent full backup.

Differential incremental backup is only supported for those workloads, where the Change Data Capture feature on the Azure SQL Server, GCP SQL Server, AWS RDS Oracle, and Azure SQL Managed Instance are enabled.

Guidelines for working with incremental backups for PaaS workloads:

- Assign a longer retention period to full backups than to incremental backups within a policy. A complete restore requires the previous full backup plus all subsequent incremental backups. It may not be possible to restore all the files if the full backup expires before the incremental backups.
- Use one storage for full and incremental backups.
- Do not create long-term copy for incremental backups.
- Do not expire random incremental backup images. Expiring them may cause application inconsistency due to data loss. NetBackup relies on the previous full backup and all the subsequent incremental backups.
- While duplicating, ensure that the full and the incremental backup copies are duplicated to the target storage. Any of the previous full or incremental images missing may result in data loss.
- While importing, ensure that the full and all the incremental backup copies are imported together. If any of the previous dependent full or incremental images are missing, it may result in failure.
- Differential incremental backup is supported on AWS RDS Oracle, only if schema changes are managed by the RDS.

About archive redo log backup for PaaS workloads

NetBackup supports archive log backup for AWS RDS Oracle workloads. Archive backups reduce the full and incremental backup windows significantly.

In this method, NetBackup backs up the data that is changed since the subsequent full or incremental backup.

Guidelines for working with archive log backups for PaaS workloads:

- It is recommended to keep the frequency of archive log backup in the protection plan under 24 hours.

- While creating a policy, assign a longer retention period to full or incremental backups than to archive log backups. A complete restore requires the previous full backup, all the subsequent incremental backups, and all subsequent archive log backups. If the full backup expires before any of the subsequent backups, it may not be possible to restore all the files.
- Use a single storage for full, incremental, and archive backups.
- Do not expire random archive backup images. Expiring them may cause application inconsistency due to data loss. NetBackup relies on the previous full backup and all the subsequent incremental and archive backups for successful restore.

Limitations and considerations

Consider the following when protecting PaaS workloads.

For all databases

- NetBackup Snapshot Manager deployment in RHEL 7.x is not supported for DBPaaS assets protection.
- NetBackup deployments in Flex Appliance and Flex Scale do not support PaaS workloads.
- Backup and restore are not supported for the databases, which makes it mandatory to use client certificates for their connection to NetBackup.
- Except for the AWS RDS workload instances, all other workload instances support only default ports; any custom ports are not supported.
- Database names containing the characters '#' and '/' are not supported for backup and restore operations. Also, the database name should adhere to the naming conventions suggested by the cloud vendors.
- ";" is not supported in server or database passwords.
- Backup and restore of a database with non-7-bit ASCII characters are not supported for a primary server running Windows or having a media server version prior to 10.1.1.
- You can duplicate the PaaS backup image to a supported storage server. But before you start a restore, you need to duplicate the image back to an MSDP server with universal share enabled. See [“Recovering duplicate images from AdvancedDisk”](#) on page 108.
- With NetBackup 10.3, you can perform backup and restore of supported Azure PaaS databases with Managed Identity database authentication. This is not supported for the Azure Database for MariaDB server. This feature requires at least one media server with version 10.2 or higher.

- For authentication of Azure databases, it is recommended to use User Assigned managed identity to work across all media servers. A database user with a system-assigned managed identity, which is associated with the media server or vm-scale-set (AKS/EKS), does not work with any other media server or media in any other vm-scale set (AKS/EKS).
- Azure Managed Identity is not supported across subscriptions of different and same tenants.
- For PaaS assets, the recovery logs are not available under **Recovery > JobID > Logs**. You can view the recovery logs either from the activity monitor or from the Restore activity tab, under asset details.
- Restore operation of the PaaS assets requires view permission for the storage server. If the storage server version is older than 10.2, additional view and create permissions for Ushare are required, along with view permissions of the storage server.
 If the logged-on user does not have view permissions for the storage server, then NetBackup tries to fetch the existing UShares during a restore. If no Ushares exist, NetBackup creates a new one named `dbpaasrestore` during the restore. NetBackup starts the recovery job subsequently.

For PostgreSQL

- Restoring security privileges is not supported.
- During restore, you can use the `-no-owner` and `-no-privileges` options. After restore, the metadata captured at the time of backup is shown as the owner/ACL in the progress log restore activity on the web UI.
- Restore does not fail if the owner or role does not exist on the destination.
- Post restore, the database has the role associated with it, according to the credentials provided in NetBackup against the destination instance.
- Users need to modify the ownership of databases post-restore.
- Azure Postgres database restore from a single to a flexible server or vice versa is not supported because of the cloud provider's limitations.
- The following characters are not supported in the database name in the restore workflow: ```, `@`, `\`, `[`, `]`, `!`, `#`, `%`, `^`, `.`, `,`, `&`, `*`, `(`, `)`, `<`, `>`, `?`, `/`, `|`, `}`, `~`, `:`, `'`, `"`, `;`, `+`, `=` and `-`.
- Uppercase username is not supported for new users added after PostgreSQL server creation.
- (RDS and Azure PostgreSQL only) SCRAM authentication configured on a database instance is not supported.

For AWS RDS PostgreSQL and AWS Aurora PostgreSQL

- Backup and restore need a media server with NetBackup version 10.4 or above and a local LSU.
- If a backup image contains a materialized view, after the restore, you need to manually refresh the materialized view. See this article to refresh materialized views:
https://www.veritas.com/support/en_US/article.100062910
- If the user credentials used to restore are those of an IAM user, then the database object gets restored with the same ownership as in the source database.
- Ownership and privileges on an object are not restored. The user that you use for the restore, becomes the owner of all the restored database objects in the following scenarios:
 - If the restored user credentials are username and password.
 - If the backup image is taken by a version before NetBackup version 10.4.

For AWS DynamoDB

- Alternate restores for regions and accounts are not supported.
- Restoring from imported images from a different primary server is only supported using the NetBackup REST API.
- The import from S3 feature during restore is supported with media version 10.3.1 or higher. This feature allows faster restores without consuming the write capacity of the tables.
- The import from S3 feature does not support the restoration of the local secondary index. This feature is enabled by default.
- To restore the local secondary index, select the option **Include local secondary indexes**. This consumes the write capacity of the tables, and the restore may take longer time.

For AWS RDS SQL

- Only the Express and Web editions for AWS RDS SQL are supported.
- For credential validation, IAM is not supported for AWS RDS SQL. You can use the username and password method.
- Only the **Amazon RDS** data management type is supported. The data management type **RDS Custom** is not supported for AWS RDS SQL instance editions.

- Databases using Transparent Data Encryption (TDE) are backed up without TDE, but using MSDP encryption. This allows for restoring your database in more scenarios, like the loss of the TDE encryption key, a cloud region outage, disaster recovery to another cloud, and so on.

For MySQL

- The restore operation requires superuser privileges if the dump file contains the CREATE DEFINER statement for backups taken on a version lower than 10.2.
- Backups taken on version 10.3 or higher cannot be restored using a version lower than 10.2.
- Backup and restore are not supported if the only SSL connection is enforced at the server-level for the GCP MySQL workload.
- You can restore a MySQL database to an alternate instance with another MySQL version than the backup instance, depending on MySQL's version compatibility.

For GCP SQL Server

- Backup and restore of read-only databases are not supported.
- Provider credentials are validated for full backup and restore and not as database credentials.
- Backup and restore of single-user-mode databases are not supported.
- If one operation is in progress, the subsequent jobs wait in the queue. If the job in progress takes time to complete, the jobs in the queue may get timed out, and fail.

For incremental backups using the GCP SQL Server

- Incremental backups after any DML changes, might fail when a table is renamed after CDC is enabled on the table. As a workaround, you must manually modify any objects that reference the renamed table. For example, if you rename a table that is referenced in a trigger, you must modify the trigger to contain the new table name. Refer to this [Azure documentation](#) link to list dependencies on the table before renaming it.
- Backup and restore of databases with binary or image data are not supported. Bulk inserts on the Cloud SQL Server requires sysadmin permission that GCP does not allow.
- While duplicating incremental backups on the different storage servers, NetBackup generates different copy numbers for the same recovery point. If you try to restore an incremental copy where no earlier full and other incremental backups are present, the restore may fail.

- If you have multiple media servers, the incremental backups can run only on version 10.3 or later.
- System databases and CDC schema are backed up and restored on the target database.
- You must set the CDC retention period greater than the period used to schedule incremental backup frequency.
- Incremental backups for databases with multiple tables can take longer to back up, as CDC enablement for multiple tables takes longer time.
- Incremental backups are not supported for database editions: Web and Express.
- Any attempts to enable CDC fail if a custom schema or a user named CDC already exists in the database.
- To ensure application consistency, NetBackup relies on the previous full backup and all the subsequent incremental backups. If a random backup image is expired, it may cause application inconsistency due to data loss.
- CDC requires SQL Server Standard or Enterprise editions. If a database is attached or restored with the KEEP_CDC option to any edition other than Standard or Enterprise, backup fails. The error message 932 is displayed.

For Azure SQL and SQL Managed Instance

- An Azure VM, which is used as a media server, should be in the same Vnet as that of the Azure-managed instance. Alternatively, if the media server and SQL-managed instance are in different Vnet, then both the Vnets must peer to access the database instance.
- Backup fails when Readlock is placed on the database or resource group.
- Backup is partially successful when the Delete lock is placed on the database or resource group. The tempdb stale entry does not get deleted from the Azure cloud portal. You need to manually delete it.
- To restore a database on an Azure SQL server or Azure Managed Instance, you must assign AAD admin privilege on the target server. Before the restore, assign the rights to any of these:
 - The system or the user-managed identity of the media servers.
 - The `vm-scale-set` in which the NetBackup media is deployed; in case of AKS or EKS deployment.

For Azure SQL Server and SQL Managed Instance incremental backup

- You can enable Change Data Capture (CDC) only on databases tiers S3 and above. Sub-core (Basic, S0, S1, S2) Azure SQL Server and SQL Managed Instance databases are not supported for CDC.
- You may encounter backup or restore issues for databases with encrypted columns in the table. As a workaround, Microsoft suggests using the Publish/Extract commands to tackle this issue.
- Restoration may fail for a database with blob data in the table.
- To duplicate incremental backups on different storage servers, NetBackup generates different copy numbers for the same recovery point. If you try to restore an incremental copy where no earlier reference of a full or other incremental backup is present, the restore fails.

Note: Incremental backups of Azure SQL Server can run only on NetBackup media server version 10.2 and above. Incremental backups of Azure SQL Managed Instances can run only on NetBackup media server version 10.3 and above.

- The user ID used for the cloud service must have permission to enable and disable CDC. Without this permission, you can see errors as follows:

```
3842: "Failed to enable CDC"
and
3844: "Failed to disable CDC"
```
- Any attempt to enable CDC fails if a custom schema or a user with the name `cdc` exists in the database. The term `cdc` is reserved for system use.
- In a database with the CDC schema created before taking the first full backup, the schema does not get backed up or restored.
- If you restore to any edition other than Standard or Enterprise, the operation is blocked because CDC requires SQL Server Standard or Enterprise editions. Error message 932 is displayed.
- Avoid backing up databases with BLOB data tables. If a table contains BLOB data, then the backup might be successful, but the restore fails.
- Encryption setting of an Azure SQL Server or Azure SQL Managed Instance database may not be preserved (*is_encryption=0*) during a restore.

For Azure Cosmos DB for MongoDB

- Discovery, protection, and restore are not supported if the account is configured using the vCore cluster.
- Backup and restore are not supported if the account is configured with a customized key.
- NetBackup does not support Azure Cosmos DB for MongoDB version 3.2.
- The **Overwrite existing database** option is not supported.
- Rules for naming databases:
 - The length of the database names must be between 3 and 63 characters.
 - Database names support all characters except #, /, ?, &, <, >, =, }, \$, {,], [, ", ', .., \.

For Azure Cosmos DB for NoSQL

- Backup and restore are not supported if the account is configured with a customized key.
- Protection of Azure Cosmos DB for MongoDB version 3.2 is not supported.
- The **Overwrite existing database** option is not supported.
- Rules for naming databases:
 - The length of the database names must be between 3 and 63 characters.
 - Database names support all characters except #, /, ?, &, <, >, =, }, \$, {,], [, ", ', .., \.

For Amazon RDS for Oracle

- Backup and restore support for full, differential incremental, and archive redo log type protection.
- Oracle 21c and 19c CDB are supported. The 19c non-CDB version is also supported.
- CDB databases with multi-tenant and single-tenant container databases, and non-CDB databases are supported.
- Oracle Enterprise and Standard Editions are supported.
- Backup and restore are both supported for S3 as the staging path.
- Backup and restore are not supported for TDE-enabled RDS Oracle instances or read replicas.
- For credential validation, IAM is not supported for AWS RDS Oracle. You can use the username and password method.

- The option group attached to RDS Oracle must have the same database engine version and the same database engine name.
- Restore is supported using the S3 staging path only, including manual recovery from the Instant access database tab.
- Only the Amazon RDS data management type is supported. The data management type RDS Custom is not supported.
- If the S3 integration is not configured or the S3 configuration fails, a backup falls back to EFS in the 19c version only, provided EFS is configured. Ensure that you remove the EFS ID entry from the option group before deleting that EFS.
- Before running the archive log backup, set the retention period in the protection plan. See the Knowledge-base article: https://www.veritas.com/support/en_US/article.100059038
- Do not take external backups using the RDS `rman` APIs on the instance to maintain data consistency.
- The recovery script supports EC2 or on-premises VMs.
- NetBackup enforces full backups in these three cases:
 - If a backup cancellation or failure occurs. NetBackup tracks such events by keeping a flag in the previous DBPaaS statefile.
 - If you schedule the first backup as an incremental or archive log backup.
 - If you take multiple incremental or archive backups, beyond the threshold value. The threshold value refers to the number of recovery points of incremental and archive log backups.

For Amazon Redshift

- Restores to alternate regions or alternate accounts are not supported.
- NetBackup protects the individual AWS Redshift cluster databases. Protection of the entire AWS Redshift cluster is not supported.
- Only user databases are protected. System databases are not displayed or protected.
- Restore from imported images from a different primary server is supported only using the NetBackup REST API.
- Only Redshift clusters are supported. Serverless Redshift is not supported.
- All clusters whose databases you are taking backup must be in the available state.
- Table names with double quotes and case-sensitive names are not restored.

- File count during the restore may show one file less than the total number of backed-up files.
- It is not recommended to take backups of databases with empty tables.
- NetBackup provides crash-consistent Redshift data protection. Consider the type of activity and application requirements before taking backups to determine if an application needs to checkpoint or quiesce for backup operations.

Discovering PaaS assets

NetBackup lets you discover, protect, and restore PaaS database assets. You can also 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: If you have upgraded the NetBackup Snapshot Manager (previously CloudPoint) from version 10.0 to 10.1. For all users with custom roles, the PaaS assets are marked as deleted in the **PaaS** tab. The assets do not show any recovery-point on them, instead, new assets with the same name are visible. The old assets get removed from the **PaaS** tab after the subsequent scheduled asset cleanup (default duration is 30 days). To resolve this issue, re-assign permissions of all the new assets to the existing RBAC role or create a new custom role. For more information, see *NetBackup WebUI Administrator's Guide*.

Note: If you change the Snapshot Manager cloud plug-in configuration from Azure service principal to Azure managed identity, the statuses of the previously discovered PaaS assets are displayed as deleted. NetBackup Snapshot Manager removes the deleted assets every 24 hours. If you want to perform backup or recovery before the scheduled cleanup, contact Veritas Technical Support.

To discover PaaS assets:

- 1 Add a Snapshot Manager. See [“Add a Snapshot Manager”](#) on page 12.
- 2 Add Microsoft Azure, GCP, or AWS as a provider. See [“Add a cloud provider for a Snapshot Manager”](#) on page 12.
- 3 Run a discovery. See [“Discover assets on Snapshot Manager”](#) on page 19.

After the discovery is complete, you can find all the discovered assets in the **PaaS** tab, in the **Cloud** workload.

All discovered AWS RDS assets appear in the **Applications** tab. The RDS instances support provider snapshot-based backups as well as NetBackup-managed backups.

NetBackup can manage and protect all the assets listed under the **PaaS** tab. Additionally, Azure SQL database and Azure SQL Managed database assets can also be backed up by Microsoft Azure.

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

Viewing PaaS assets

To view PaaS assets:

- 1 On the left, click **Workloads > Cloud**.
- 2 In the **PaaS** tab, the assets that are available to you are displayed. The RDS assets are displayed in the **Applications** tab.

You can perform: **Add protection**, **Backup now**, and **Manage credential** operations on the displayed assets.

For DynamoDB and Amazon Redshift assets, the **Manage credentials** option is not available.

For the deleted assets, you can only manage credentials.

Managing PaaS credentials

You can add credentials for a database listed in the **PaaS** and **Applications** tab under **Cloud** workload. You can add, edit, or delete PaaS credentials from the central **Credential management** console in NetBackup. Some workloads, like DynamoDB and Amazon Redshift, do not support credential management through NetBackup and leverage the provider credentials.

View the credential name that is applied to a database

You can view the named credential that is configured for the databases in the **Credential name** column of the **PaaS** tab. If the credentials are not configured for a particular asset, this field is blank.

To view the credentials for PaaS databases:

- 1 On the left, select **Workloads > Cloud > PaaS** tab.
- 2 Click **Show or hide columns** above the database list table.
- 3 Select the **Credential name** to display the credential name column.

Add credentials to a database

You can add or modify credentials for a database listed in the **PaaS** tab.

To add or change credentials

- 1 On the left, click **Workloads > Cloud**.
In the **PaaS** tab, the assets that are available to you are displayed. The RDS assets are displayed in the **Applications** tab.
- 2 Select the database in the table, then click **Manage credentials**.
- 3 Select a **Validation host**. The validation host must be a RHEL media server having connectivity with the PaaS workload, or a NetBackup Snapshot Manager. If you use a NetBackup Snapshot Manager, a datamover container is added to the Snapshot Manager host.

You can add existing credentials or create new credentials for the database:

- To select any existing credentials for the account, select the **Select from existing credentials** option, select the required credential from the table below, and click **Next**.
- To add a new credential for the account, select **Add credentials**, and click **Next**. Enter a **Credential name**, **Tag**, and **Description** for the new credential. Under **Service credentials**:
 - Select **Role based database authentication (Applicable for supported database service)** to use AWS IAM, Azure System Managed and User Managed authentication.
 - Select **IAM database authentication (Applicable for Amazon RDS only)** for Amazon RDS assets only, and specify a **Database user name**.

See [“Creating an IAM database username”](#) on page 84.

Note: If the Snapshot manager is deployed in the cloud with an attached IAM role having the required permission, you must also deploy the media server in the same cloud environment and attach the same IAM role. Otherwise, the backup jobs for the AWS assets fail.

Note: If the media server or NetBackup Snapshot Manager instance has instance metadata service (IMDSv2) enabled, then ensure that the *HttpPutResponseHopLimit* parameter is set to 2 on the hosting VM. If the value of the *HttpPutResponseHopLimit* parameter is not set to 2, then the AWS calls to retrieve the metadata from the media server or the NetBackup Snapshot Manager containers created on the VM fail. For more information on the IMDSv2 service, refer to [Use IMDSv2](#) in Amazon documentation.

- Select **Azure System Managed Identity authentication** or **Azure User Managed Identity authentication** as required. Enter the username of the database, and click **Next**.
 To perform backup and restore operations utilizing managed identity authentication, you must configure AAD admin to the source and target database servers.
 See [“Creating a system or user-managed identity username”](#) on page 86.

Note: If the snapshot manager is deployed in the cloud with an attached Managed Identity having the required permissions, attach the same identity to the media server. For AKS and EKS deployments attach the same Managed identity to the VM scale set.

- Select **Password authentication** and specify the username and the password for the database server.
 If you are using AWS RDS Oracle, the username must be in the format *username@tenantdatabasename* for AWS RDS Oracle Multitenant deployment architecture.
 If you are using Azure Cosmos DB for NoSQL:
 - Username is the **Account URI** that you can find on the Azure portal, at **Settings > Keys > URI**.

- Password is the **Primary Key** or **Secondary Key** that you can find on the Azure portal, at **Settings > Keys > PRIMARY KEY** or **SECONDARY KEY**.
- Read keys can only take backup. It is recommended to use read-write keys to restore databases.

If you are using Azure Cosmos DB for MongoDB:

- Username is the same as the account name that you can find on the Azure portal, at **Settings > Connection Strings > USERNAME**.
- Password is the **Primary Key** or **Secondary Key** that you can find on the Azure portal, at **Settings > Keys > PRIMARY KEY** or **SECONDARY KEY**.
- Read keys can only take backup. It is recommended to use read-write keys to restore databases.

Click **Next**.

- Add a role that you want to have access to the credential. To add new permissions to a role:
 - Click **Add**.
 - Select a role.
 - Select the credential permissions that you want the role to have.
 - Click **Save**.

4 Click **Next** to finish creating the credential.

For more information about credentials and how to edit or delete a credential, see *NetBackup Web UI Administrator's Guide*.

Creating an IAM database username

To create an IAM username:

- 1** Enable IAM DB authentication on the RDS DB instance.
- 2** Create the Database user, using the master login (rds_iam)
 - For MySQL create the username using master login (rds_iam):
 - `mysql --protocol=tcp --host=instance_fqdn --user=admin -p --port=3306`
 - `CREATE USER iamuser IDENTIFIED WITH AWSAAuthenticationPlugin as 'RDS';`
 - `GRANT SELECT, INSERT, UPDATE, DELETE, CREATE, DROP, RELOAD, PROCESS, REFERENCES, INDEX, ALTER, SHOW DATABASES, LOCK TABLES, CREATE VIEW, SHOW VIEW, CREATE`

```
ROUTINE, ALTER ROUTINE, EVENT, TRIGGER ON *.*
'<db_user>'@'%';
```

- For PostgreSQL, create the user on the server.
 - `psql -h instance_fqdn -U postgres`
 - `CREATE USER iamuser WITH LOGIN;`
 - `GRANT rds_iam TO iamuser;`
 - `ALTER ROLE iamuser WITH LOGIN CREATEDB;`
 - `GRANT rds_superuser TO iamuser;`

3 Attach the RDS policy to the IAM role attached to the NetBackup media server.

For more details, see *AWS permissions required by NetBackup Snapshot Manager* section in the latest version of the *NetBackup Snapshot Manager Install and Upgrade Guide*.

Configuring permissions for the database user

For MySQL

Create a database user with a master login and grant these permissions:

- `mysql --protocol=tcp --host=instance_fqdn --user=admin -p --port=3306`
- `CREATE USER dbuser IDENTIFIED BY '<password>';`
- `GRANT SELECT, INSERT, UPDATE, DELETE, CREATE, DROP, RELOAD, PROCESS, REFERENCES, INDEX, ALTER, SHOW DATABASES, LOCK TABLES, CREATE VIEW, SHOW VIEW, CREATE ROUTINE, ALTER ROUTINE, EVENT, TRIGGER ON *.* TO `dbuser`@'%' WITH GRANT OPTION;`

For PostgreSQL

Create a database user under the server and grant the following permissions:

- `psql -h instance_fqdn -U postgres`
- `CREATE USER dbuser WITH PASSWORD '<password>' CREATEDB;`
- (For AWS RDS PostgreSQL) `GRANT rds_superuser TO dbuser;`
- (For AZURE PostgreSQL) `GRANT azure_pg_admin TO dbuser;`
- (For GCP PostgreSQL) `GRANT cloudsqlsuperuser TO dbuser;`

For SQL Server

Create a database user under the server and grant the following permission:

- Create a login on the server:

```
CREATE LOGIN dbuser WITH PASSWORD='<password>'
```

- Create a user for the database in the server:
 - CREATE USER [dbuser] FOR LOGIN [dbuser]
 - ALTER ROLE [db_owner] ADD MEMBER [dbuser]

Note: The database user must not be part of any database deny role. For example: db_denydatareader and db_denydatawriter.

Creating a system or user-managed identity username

For Azure SQL Server and Managed Instance

Do any of the following configurations:

Configure the managed identity user as an AAD admin:

- Set the AAD admin on the SQL server or the Managed instance.
- Go to Settings > Azure Active Directory > Set admin. Search and set the system-assigned or user-assigned managed identity, and save.

Note: Only media servers configured with both system-assigned managed identity and AAD administrator permissions can perform backup and restore.

Create a managed identity user on the database using the SSMS client:

- To set AAD admin for SQL server, create users, go to Settings > Active Directory admin > Set admin. Pick active directory for the user, and save.
- Login to the SQL database or Managed database to create a user under that database.

```
CREATE USER [<managed_identity>] FROM EXTERNAL PROVIDER;
ALTER ROLE db_owner ADD MEMBER [<managed_identity>];
```

- Provide login permission for that user on the SQL Server, run

```
# CREATE USER [<managed_identity>] FROM EXTERNAL PROVIDER;
# ALTER ROLE loginmanager ADD MEMBER [<managed_identity>];
```

Note: You must create users for all media servers communicating with the database using the system-assigned managed identity.

Note: To restore a database, you must configure the managed identity user as an AAD admin on the target server.

For MySQL

- To configure the AAD admin for the MySQL server, create a user. Go to Settings > Active Directory admin > Set admin. Pick the active directory user, and save.

- Get the client ID for managed identity using Azure CLI, run

```
# az ad sp list --display-name <managed_identity> --query [*].appId  
--out tsv
```

- Generate an access token to log on, using Azure CLI, run:

```
# az account get-access-token --resource-type oss-rdbms
```

- Log on using the AAD admin user and access token, run:

```
# mysql -h <server name> --user <user name>  
--enable-cleartext-plugin --password=<token>
```

- Create the manage identity user and grant the permissions, run:

```
# SET aad_auth_validate_oids_in_tenant = OFF;  
# CREATE AADUSER '<db_user>' IDENTIFIED BY  
'<Generated_client_id>';  
# GRANT USAGE, DROP, SELECT, CREATE, SHOW VIEW, EVENT, LOCK  
TABLES , ALTER, CREATE VIEW, INSERT, REFERENCES, ALTER ROUTINE,  
PROCESS ON *.* TO '<db_user>'@'%'
```

For PostgreSQL

- To configure the AAD admin for the PostgreSQL server, create a user. Go to Settings > Active Directory admin > Set admin. Pick the active directory user, and save.

- Get the client ID for the managed identity:

```
# az ad sp list --display-name <managed_identity> --query  
[*].appId --out tsv
```

- Generate the access token required to login, run:

```
# az account get-access-token --resource-type oss-rdbms
```

- Export the password for the generated token, run:

```
# export PGPASSWORD=<token>
```

- Login using the AAD admin user and the access token, run:

```
# psql "host=<host name> port=5432 dbname=<dbname> user=<user name> sslmode=require"
```

- To create a user and grant permissions, run:

```
# SET aad_auth_validate_oids_in_tenant = OFF;
# CREATE ROLE <db_user> WITH LOGIN PASSWORD '<client_id>' IN ROLE azure_
# GRANT azure_pg_admin TO <db_user>;
# ALTER USER smipguser CREATEDB;
# ALTER USER smipguser Replication;
```

Note: Only user-managed identity is supported for MySQL Flexible Server. Managed Identity support is not available for PostgreSQL Flexible Server.

For Azure Cosmos DB for NoSQL

1. Log on to your Azure portal.
2. To assign the **Cosmos DB Built-in Data Contributor** role to the managed identity, run the command:

```
# az cosmosdb sql role assignment create -a <Account_Name> -g
<Resource_Group_Name> -s "/" -p <Object_ID/Principle_ID> -d
00000000-0000-0000-0000-000000000002
```

Where:

- *Account_Name* is the Azure Cosmos account name.
- *Resource_Group_Name* is the Resource group name of the account.
- *Object_ID/Principle_ID* is the Managed identity object or principle ID.
- *00000000-0000-0000-0000-000000000002* is the **Cosmos DB Built-in Data Contributor** role ID.

Add protection to PaaS assets

After you discover the PaaS assets, you can add protection to them in the **Applications** or **PaaS** tab in **Cloud** workload.

To add protection to PaaS assets

- 1** On the left, click **Workloads > Cloud**.
- 2** To protect AWS RDS-supported database assets, click the **Applications** tab. For other PaaS assets, click the **PaaS** tab.
- 3** Check if the asset that you want to protect has a credential.
See [“View the credential name that is applied to a database”](#) on page 82..
If the **Credential name** column is empty, you need to assign a credential to the asset.
See [“Add credentials to a database ”](#) on page 82.
- 4** To add protection to an asset, select the asset and click **Add protection**.
An asset must have assigned credentials to be eligible for most operations. For example, if you want to assign the asset to a protection plan or perform backup now.
- 5** Select a protection plan and click **Next**.
- 6** Review the configuration settings and click **Protect**.

Perform backup now

Using this option you can create a one-time backup of the selected asset. This backup does not affect any future, or scheduled backups.

To perform a backup now

- 1 On the left, click **Workloads > Cloud**.

To back up AWS RDS supported database assets, click the **Applications** tab.
For other PaaS assets, click the **PaaS** tab.

Note: You can see and protect the user-created databases. The system databases are not shown and protected, as these databases need the cloud provider's superuser privilege to perform backup and restore.

- 2 Select the asset, then click **Add protection**.
- 3 Select the required protection plan, then click **Start backup**.

You can view the status of the backup job in the Activity monitor.

The database agents access the database from within the media server (container, in case of NetBackup deployed in AKS and EKS environments), and perform NFS mount of the Universal share path on the media server (backup host).

Note: For incremental backup of Azure SQL databases, GCP SQL Server, and AWS RDS Oracle, NetBackup performs a full backup even if a protection plan with backup type differential incremental protects the asset.

Note: For archive redo log type schedules, NetBackup takes full backup for backup now, irrespective of what you specified in the protection plan.

Recovering cloud assets

This chapter includes the following topics:

- [Recovering cloud assets](#)
- [Perform rollback recovery of cloud assets](#)
- [Recovering AWS or Azure VMs to VMware](#)
- [Recovering PaaS assets](#)

Recovering cloud assets

You can restore AWS, Azure, Azure Stack, OCI, and GCP VM assets from snapshot copy, replica copy, backup copy, or duplicate copy. You can also restore backup images from AWS EC2 or Azure VMs to on-premises VMware VMs.

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 zone, to a different subscription, and restore VMs or disks to a different resource group.

- For GCP: Select **Firewall rule**.
- For Azure: Select **Network security group**.
- For AWS: Select **Security group**.
- For OCI: Select **Network security group**.

About the pre-recovery check for VMs

The 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 of 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 Snapshot Manager and cloud credential validation.
- Validity of selected encryption keys.

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 2-1 Supported parameters for Azure, Azure Stack, GCP, and AWS snapshot and backup copies

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

Table 2-1 Supported parameters for Azure, Azure Stack, GCP, and AWS snapshot and backup copies (*continued*)

Change provider configuration					Y	Y		
Change resource group for disks	Y	Y			Y	Y		
Zone/Availability domain	Y		Y	Y	Y		Y	Y
Security group/Firewall rule/Network security group	Y	Y	Y	Y	Y	Y	Y	Y
Edit disk encryption	Y		Y	Y	Y		Y	Y

Recovering virtual machines

To recover a VM

- 1 On the left, click **Workloads > 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 to 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 replicas 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, and 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 and OCI 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**.

During VM restore from snapshot or backup copy, encryption keys can be selected from individual disks or all disks at the same time as follows:

- Select the **Volume** and click **Edit the encryption key** option.
- Select the required **Encryption type**.
- Select the required encryption **Key** and click **Save**.

8 Click **Next**.

9 In the Recovery options page:

- If you restore a backup copy, to restore to a different zone, select a **Zone** or **Availability domain**. To select a network available in that zone, click the search icon near **Network configuration**, and select a target network for recovery.

You can also select:

- For GCP: **Firewall rule**
- For Azure: **Network security group**
- For AWS: **Security group**
- For OCI: **Network security group**
- (*Only for GCP*) If you restore a snapshot copy, to restore to a different region, select a **Region**. To select a network available in that zone, click the search icon in **Network configuration**, and select a target network for recovery. The list shows networks available in that zone.

- If you restore a snapshot copy, to restore to a different zone, select a **Zone** or **Availability domain**. To select a network available in that zone, click the search icon in **Network configuration**, and select a target network for recovery. The list shows networks available in that zone or Availability domain.
You can also select **Security group / Network security group / Firewall rule** for AWS, Azure, OCI, and GCP cloud providers respectively.

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 a 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 any 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.

Note: Selection of **None** for zone means that the VM is not placed in any zone and selection of **None** for **Network security group/Security group/Firewall rule** means that no security rules are applied to the restored VM.

- 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 their 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 **Workloads > 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 to recover.
- 4 Click the **Recovery points** tab. In the calendar view, click the date on which the backup occurred.
The available images are listed in rows with the backup timestamp for each image.
- 5 On the 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 the 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 **Workloads > 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 to recover.
- 4 Click the **Recovery points** tab. In the calendar view, click the date on which the backup occurred.
The available images are listed in rows with the backup timestamp for each image.
- 5 On the 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.

Note: (Applicable for Azure cloud) Application restore to an alternate location for the ADE-enabled VMs is not supported.

Recovery scenarios for GCP VMs with read-only volumes

The following table describes how NetBackup handles the restore/recovery of GCP VMs that have read-only volumes.

Table 2-2 Recovery scenarios for read-only GCP VMs

Scenario	Handling
Restoring a volume from the snapshot of an attached read-only disk, from the Volumes tab under Cloud workload.	During restore, the disk is 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 Virtual machines tab under Cloud workloads.	During the restore of such a VM to its original or alternate location, a read-only disk is restored in a read/write mode.
Restoring a VM with a read-only disk, from an app-consistent snapshot, from the Virtual machine tab under Cloud workload.	<p>You can attach a read-only disk to multiple VMs, but NetBackup discovers it under only one VM.</p> <p>For a Windows VM, the snapshot fails with a VSS error, similar to the following:</p> <p>Failure: flexsnap.GenericError: Failed to take snapshot (error: Failed to create VSS snapshot of the selected volumes.)"</p> <p>For a Linux VM, the snapshot may or may not be successful for the VM under which the disk is discovered, but fails for the rest of the VMs due to the missing dependencies. Error example:</p> <p>linear_flow.Flow: create snapshot (test-win) of host linux-1(len=4) requires ['snap_google-gcepd-us-west 2-b-7534340043 132122994'] but no other entity produces said requirements\nMissingDependencies</p> <p>In the above case, if a snapshot is successful for a Linux VM, a read-only disk is restored in a read/write mode.</p>

When taking a snapshot or backup from snapshot of a source VM, additional information about disks is saved. The **autoDelete** flag determines whether to delete

the disk when deleting the VM. Hence, if a new VM is created from snapshot or backup from snapshot, then disks are set as the source VM.

For example,

Source VM:

Disk1: **autoDelete** is set to true (When the source VM is deleted and **autoDelete** is set to **true** then the disk is deleted automatically)

Disk2: **autoDelete** is set to false.

Restored VM:

Disk1_suffix: **autoDelete** is set to true.

Disk2_suffix: **autoDelete** is set to false.

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, OCI, and GCP workloads do not support rollback restore.

To perform rollback recovery of the cloud asset

- 1 On the left, click **Workloads > 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 AWS or Azure VMs to VMware

NetBackup lets you restore cloud-based backup images from AWS EC2 or Azure VMs to on-premises VMware VMs.

Pre-requisites

- The recovery host must run on the RHEL platform. For the recovery host version, see the *VMware Compatibility* section in *Enterprise Server and Server 10.0 - 10.x.x OS Software Compatibility List*.
- When recovering VM, the supported transport mode for VMware is NDB.
- It is recommended that you use a different VM name for the converted VM, which is not used by any existing VMs in the VMware server.

To recover a cloud VM to VMware:

- 1 On the left, click **Workloads > Cloud**.
- 2 Click the **Virtual Machines** tab.
- 3 Double-click the protected asset that you want to recover, and then click the **Recovery points** tab.

The available images are listed in rows with the backup timestamp for each image.
- 4 In the **Copies** column, click the copy that you want to recover. You can recover only backup images.
- 5 Click the ellipsis menu (three dots) in the row of the copy, and click **Restore Virtual Machine**.
- 6 In the Recovery target page, do the following:
 - Select **Provider** as **VMware**.
 - **Display name**: To change the display name, enter a new one in the field.
 - **ESXi server or cluster**: Select the ESXi server or the cluster where the VM resides.
 - **Folder**: Specify the folder that contains the VM.
 - **Resource pool or vApp**: Specify the resource pool for the VM.
 - **Datastore or datastore cluster**: Specify the datastore for the VM and the disks.
 - **Network configuration**: Select the network switch on the ESXi server.
- 7 Click Next.
- 8 In the Recovery options page, do the following:
 - **Recovery host**: Select the host that you want to use to perform the recovery.

- **Power on after recovery:** (Optional) Select to keep the VM powered on after recovery.
 - **CPU number:** Specify the CPU number of the converted VM.
 - **Memory size (GB):** Specify the memory size of the converted VM.
- 9 Click Next. Review the parameters, and click **Start recovery**.
- The Restore activity tab shows the job progress.

Post-recovery considerations for cloud VMs recovered to VMware

Considerations for the restored VMs:

- These are some default Configurations of the recovered VMs. You may need to modify them manually before using the VMs.
 - The default boot configuration is EFI.
 - the default disk controller is SCSI.
 - Default Network adapter type is VMXNET 3.
 - Default OS type is x64.
- In Azure, for generation 1 VMs, you must modify the boot configuration in "VM options" to BIOS. Also, change the disk controller to "IDE". For more information, see Azure documentation:

<https://docs.azure.cn/en-us/virtual-machines/generation-2?view=azs-2102>

See “[Protecting AWS or Azure VMs for recovering to VMware](#)” on page 57. for the pre-requisites and considerations for the backed-up cloud VMs.

Steps to recover images from cloud VMs to VMware

This section contains the outline steps you need to follow to recover different types of cloud VMs to VMware. Ensure that you have performed a full backup of the source VM to an MSDP storage server. See “[Protecting cloud assets or intelligent groups for cloud assets](#)” on page 26.

See “[Protecting AWS or Azure VMs for recovering to VMware](#)” on page 57. for the detailed pre-requisites and considerations for the backed up cloud VMs.

Recovering images from AWS to VMware

Windows Server 2022

Prerequisites for the backed-up cloud images:

- Change the network interface to use DHCP, enabled on boot.

- Create a local administrator before backup.

To recover Windows 2022 VM images to VMware:

- 1 Recover the image using NetBackup. See [“Recovering AWS or Azure VMs to VMware”](#) on page 98.
- 2 Log on to your VMware server, and edit the converted VM settings. In the VM Options page, click Boot Options, then change Firmware to BIOS.
- 3 Obtain the IP address to log on the converted VM through RDP.

RHEL 9.x**Prerequisites for the backed-up cloud images:**

- Change the network interface to use DHCP, enabled on boot.
- Create a new user to log on to the recovered VM.

To recover RHEL 9.x VM images to VMware:

- 1 Recover the image using NetBackup. See [“Recovering AWS or Azure VMs to VMware”](#) on page 98.
- 2 Log on to your VMware server, and edit the converted VM settings. In the VM Options page, click Boot Options, then change Firmware to BIOS.
- 3 Obtain the IP address to log on to the converted VM through SSH.

SUSE 15SP5**Prerequisites for the backed-up cloud images:**

- Change the network interface to use DHCP, enabled on boot.
- Create a new user to log on to the recovered VM.

To recover SUSE 15SP5 VM images to VMware

- 1 Recover the image using NetBackup. See [“Recovering AWS or Azure VMs to VMware”](#) on page 98.
- 2 Obtain the IP address to log on to the converted VM through SSH.

Recovering images from Azure to VMware**Windows 2022****Pre-requisites for the backed-up cloud images:**

- Change the network interface to use DHCP, enabled on boot.

To recover Windows 2022 VM images to VMware:

- 1 Recover the image using NetBackup. See [“Recovering AWS or Azure VMs to VMware”](#) on page 98.
- 2 For Windows 2022 Gen 1, log on to the VMware server, and edit the converted VM settings. In the VM Options page, click Boot Options, and then change Firmware to BIOS.
- 3 Obtain the IP address to log on the converted VM through RDP.

RHEL 9.x**Pre-requisites for the backed-up cloud images:**

- You need the VMW_PVSCSI driver in the source VM. To see if the driver already exists, run:

```
lsinitrd | grep -i vmw_pvscsi
```

To install the driver, do the following:

- To backup `initramfs`, run the following commands one by one:

```
cd /boot
```

```
cp initramfs-`uname -r`.img initramfs-`uname -r`.img.bak
```

- To open the `dracut.conf` file, run:

```
vi /etc/dracut.conf
```

Uncomment the line `#add_drivers+=`. Add the value `"vmw_pvscsi"` to the line, separating the existing module from the space. For example:

```
# additional kernel modules to the default.  
add_drivers+="vmw_pvscsi"
```

- To create new initial ramdisk images with new modules, run:

```
dracut -f -v -N
```

- Run any of the following commands to check if the new modules exist in new initial ramdisk images:

```
lsinitrd | grep -i vmw_pvscsi
```

```
lsinitrd -f /boot/initramfs-`uname -r`.img | grep -i vmw_pvscsi
```

- Create a new user to log on to the recovered VM.
- Change the network interface to use DHCP, enabled on boot.

To recover RHEL 9.x VM images to VMware:

- 1 Recover the image using NetBackup. See [“Recovering AWS or Azure VMs to VMware”](#) on page 98.
- 2 For RHEL Gen 1, log on to the VMware server, and edit the converted VM settings. In the VM Options page, click Boot Options, and then change Firmware to BIOS.
- 3 Obtain the IP address to log on the converted VM through SSH.

SUSE 15SP5**Pre-requisites for the backed-up cloud images:**

- Change the network interface to use DHCP, enabled on boot.
- Create a new user to log on to the recovered VM.

To recover SUSE 15SP5 VM images to VMware

- 1 Recover the image using NetBackup. See [“Recovering AWS or Azure VMs to VMware”](#) on page 98.
- 2 If you do not have an existing vmw_pvscsi driver in the source VM, log on to the VMware server, and edit the Converted VM settings. In the Virtual Hardware page, click Hard disk, and then change the Virtual Device Node to IDE.
- 3 For SUSE 15SP5 Gen 1, log on to the VMware server, and edit the converted VM settings. In the VM Options page, click Boot Options, and then change Firmware to BIOS.
- 4 Obtain the IP address to log on to the converted VM through SSH.

Recovering PaaS assets

PaaS assets are listed under the **Cloud** workload. You can restore Amazon RDS assets from the **Applications** tab. All other PaaS assets are available for restore from the **PaaS** tab. Recovery flows for Azure assets are different, based on whether they are NetBackup-protected or Azure-protected.

In NetBackup 10.3 and later, you can separately restore the data or schema and the metadata for the MySQL database. You need superuser privileges to restore the metadata and at least one media server at version 10.2 or later.

Note: For a MySQL restore, if you do not have admin or root user privileges, then you must have the view permission, along with the restore permissions.

PaaS assets support instant access during recovery. Instant access enables faster access to data and reduces overall recovery time.

Note: While viewing PaaS restore jobs in the Activity monitor, the fields **Bytes transferred** and **Estimated bytes remaining**, may not indicate correct information. You can look at the number of **Files written** for the correct status, and the NetBackup logs.

Recovering non-RDS PaaS assets

You can restore the non-RDS PaaS assets from the **PaaS** tab, under Cloud workload.

To restore non-RDS PaaS assets:

- 1 On the left, click **Workloads > Cloud** and click the **PaaS** tab. Click the name of the asset that you want to recover.
- 2 Click the **Recovery points** tab, for Azure assets, additionally select **NetBackup managed**.

The available recovery points are displayed in the table.

- 3 Click **Recover** in the row of the image that you want to recover.
- 4 In the **Name** field, the original name of the asset appears by default. You can change the name in the field. You may not be able to change this name later.
- 5 (Optional) In the **Target instance** field, the source instance of the asset is selected by default. To restore to an alternate instance, select the required instance. **Target instance** is not available for DynamoDB assets.
- 6 (Optional, for MySQL databases only.) Select **Restore metadata** to restore metadata such as views, triggers, store procedures, and so on.
- 7 (Optional, for MySQL databases only.) For the target instance credentials for restore:
 - Select **Use already associated credentials** to use the credentials that are already associated with the instance, and click **Start recovery**.
 - Select **Use different credentials** to use a different set of credentials, either existing credentials or create a new one.
See [“Add credentials to a database”](#) on page 82.
The validation host for validating these credentials must be the same as the one used during backup. If the host used during backup is not available during credential validation during restore, then validation fails.

(Optional) Select **Make default credentials** to set these credentials as default credentials for the asset.

8 Click **Start recovery**.

The **Restore activity** tab shows you the status.

Recovering RDS-based PaaS asset

You can restore the RDS-based PaaS assets from the **Applications** tab, under the **Cloud** workload.

To restore RDS-based PaaS assets:

- 1** On the left, click **Workloads > Cloud** and click the **Applications** tab. Click the name of the asset that you want to recover.
- 2** Click the **Recovery points** tab in the calendar, select the date for which you want to see the recovery points.

The available recovery points are displayed on the right.

- 3** Click **Recover** in the row of the image that you want to recover.
- 4** Under **Source databases**, select the databases that you want to restore. Click **Add database**, in the **Add database** dialog, select the required databases, and click **Select**.
- 5** (For Amazon RDS Oracle databases only) Enter the staging path in the **AWS S3 bucket name** field. Click **Start recovery**. The recovered database appears in the **Instant access databases** tab. Recovery is possible on self-managed instance EC2 or on-premises VM To complete the recovery of the asset, see the Knowledge base article:

https://www.veritas.com/support/en_US/article.100058945

You can select a different S3 bucket to stage the restored data, than the one used during the backup. You can also select a S3 bucket at a different region.

- 6** Enter a prefix to add to the restored databases, or use the default. This field must have a value.
- 7** (Optional) In the **Target instance** field, the source instance of the asset is selected by default. To restore to an alternate instance, select the required instance.
- 8** (Optional, for MySQL databases only.) Select **Restore metadata** to restore metadata such as views, triggers, store procedures, and so on.
- 9** (Optional, for MySQL databases only.) For the target instance credentials for restore:

- Select **Use already associated credentials** to use the credentials that are already associated with the instance, and click **Start recovery**.
- Select **Use different credentials** to use a different set of credentials, either existing credentials or create a new one.
See [“Add credentials to a database”](#) on page 82.
(Optional) Select **Make default credentials** to set these credentials as default credentials for the asset.
- Select a validation host to validate the provided credentials.

10 Click **Start recovery**.

The **Restore activity** tab shows you the status.

These two restore workflows implicitly create an instant access mount share against the recovery point.

Recovering Azure-protected assets

NetBackup lets you restore Azure SQL database and Azure SQL managed database assets that are 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 is 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 **Workload > Cloud**.
- 2** Click the **PaaS** tab.
All the discovered PaaS assets are displayed.
- 3** Under **Recovery points type**, select **Provider protected**.
- 4** Click **Restore** in the row of the protected Azure SQL database and Azure SQL managed database asset that you want to recover.
- 5** In the **Recovery points** tab, under **Point in time backup**, click **Restore**.
- 6** 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.

The default restore date and time displayed in web UI 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 databases, the default restore time is 6 minutes earlier than the current time.

- 7 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.

- 8 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.

- 9 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 **Workloads > 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 to 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 *restore_<dbName>* 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 the portal as well as Snapshot Manager 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 runs.

Recovering duplicate images from AdvancedDisk

A 10.1 media server cannot initiate PaaS restores from a duplicated image, if the image resides on an AdvancedDisk storage or an MSDP cloud storage. As a workaround, you can perform the following steps:

Pre-requisite:

1. For AdvancedDisk the media server version associated with the MSDP server must be 10.1 or above.
2. For MSDP cloud storage, the media server version used for recovery must be 10.1.1.
3. Ensure that the ushare is set up and configured on the MSDP server.
4. Create a universal share on this MSDP storage server. Ensure that you add the corresponding media server host name/IP in the export list of ushare.

To recover from AdvancedDisk, do the following:

- 1 Using the Catalog in the web UI, manually duplicate the image to an MSDP storage. See *NetBackup WebUI Administrator's Guide* for details.

Note: To duplicate from a second copy, click search again after selecting the duplicate option in the catalog view.

- 2 Once the duplication job is completed, ensure that the new recovery point is visible for the given asset in the web UI.

To start a restore job, See [“Recovering PaaS assets”](#) on page 103.

To restore using REST API, see section:

`recovery/workloads/cloud/scenarios/asset/recover`. Refer to NetBackup API documentation.

Note: For RDS instance recovery, NetBackup does not display any error or warning messages, if you initiate the restore from a backup image residing on AdvancedDisk storage.

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 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 or backup 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 or backup using an indexing process. You can restore specific files from a snapshot only if

that snapshot has been indexed by NetBackup. You can also restore specific files from a backup only if NetBackup has indexed the backup.

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

Table 3-1 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 > Snapshot Managers > Snapshot Manager > 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 a backup job and indexing or use the Backup now option. The backup job starts immediately.
<ul style="list-style-type: none">■ Restore file or folder■ Restore volumes Note: Restore volumes is not supported for backup copy.	Perform granular restoration of a file, folder, or volume.

Supported environment list

The following table lists the supported versions.

Table 3-2 Supported versions

Application	Version
NetBackup	10.4

Table 3-2 Supported versions (*continued*)

Application	Version
NetBackup backup host OS	RHEL 7.x and 8.8
Snapshot Manager host OS	<ul style="list-style-type: none">■ RHEL 7.x and later, RHEL 8.6■ Ubuntu 18.04 LTS and 20.04 LTS Note: The version of the OS (Ubuntu 20.04 LTS) listed on the UI is the version of the container.
Cloud providers	<ul style="list-style-type: none">■ Amazon Web Services■ Microsoft Azure■ Microsoft Azure Stack Hub■ Google Cloud Platform
Snapshot Manager 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
Snapshot Manager agent host to be protected	<ul style="list-style-type: none">■ Linux OS: RHEL 7.x, and RHEL 8.8■ Windows OS Version: 2012 R2, 2016, 2019, and 2022

List of supported file systems

The following table provides details about supported file systems.

Platform	Discovered file system	Partition layouts
RHEL (With consistent snapshot property)	<ul style="list-style-type: none">■ ext3■ ext4	<ul style="list-style-type: none">■ GPT■ MBR
Note: For GCP, if the agent host is on operating system version RHEL 8.x, then Snapshot Manager must be installed on the host having operating system version RHEL 8.x.	<ul style="list-style-type: none">■ xfs	<ul style="list-style-type: none">■ No layout (direct FS)

Platform	Discovered file system	Partition layouts
Windows (With consistent snapshot property)	NTFS	<ul style="list-style-type: none">■ GPT■ MBR

Note: Application consistent snapshot is not supported for the 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 Snapshot Manager and VM to be protected with granular restore enabled have the following requirements:

- The following requirements apply to snapshots:
 - (Microsoft Azure and Azure Stack Hub) Even if Snapshot Manager 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 the snapshot-only protection plan schedule, for both Azure and Azure Stack Hub, you need to deploy the Snapshot Manager host in the same subscription and region as the VMs.
 - (Amazon AWS and Google Cloud Platform): The Snapshot Manager host and the connected VM must be in the same account and region.
 - The cloud plug-in must be configured to protect the assets in the region in which the Snapshot Manager host is deployed.
- The host must be in a connected state and must have the 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 must be attached to the Snapshot Manager instance explicitly.
- File systems on the host must be supported.

See [“List of supported file systems”](#) on page 112.

- Configure ports 5671 and 443 for open Snapshot Manager host.
- For agentless restore, in Linux systems, configure port 22 on the indexable virtual machines. For the Windows platform, configure the ports 135, 445, and the dynamic or fixed WMI-IN port on the indexable virtual machines.
- Ensure that you have appropriate permissions to perform a granular restore. See the information on role permissions in the *NetBackup Web UI Administrator's Guide*.
- Ensure that the following points are addressed before you perform the single file restore from a snapshot backup:
 - You have NetBackup and Snapshot Manager version 10.2 or later.
 - A granular restore is successful only if the backup image is restored from the MSDP storage server (10.3 or later), with instant access enabled.
 - For MSI and RPM-based agent installation, the target host agent must be upgraded to the latest version.
 - On the Windows target host, the administrator must have the attach and detach policy enabled for the disks. For more information, refer to the [AttachVirtualDisk function](#).
 - (For Windows) To restore symlink, the agent must be configured using the required access. For this, add the administrator user in the **Create symbolic links** policy under `Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment`.
 - The backup must be taken with the **Granular File and Restore** option selected.
 - The target virtual machine must have access to the MSDP storage server over NFS/SMB.
 - (For Linux) The NFS client (`nfs-utils`) must be installed to restore over NFS.
 - If `/etc/hosts` entry is created for the MSDP host in the MSDP storage server, add the FQDN of the MSDP storage server also in the same entry.
 - The Windows target must meet the following requirements:
 - (For restoring Windows image content with Restoring Access Control list) The Samba user credentials must be stored in the Windows credential manager for an MSDP storage server. This server is the one that exports the instant access share.

On the MSDP server, run the following command to generate the Samba credentials.

```
smbpasswd -a <username>
```

Add the DNS name or the IP address of the MSDP server. Provide the username from the previous step and the password that was generated in the Windows credential manager.

The `smbpasswd` command fails if the username is not present on the MSDP server. You must first add users with the command `useradd <username>` command.

- (For restoring Linux image content) The NFS client is installed.

For more information on how to enable SMB/IA on MSDP, refer to the *NetBackup Deduplication Guide*.

Verify the SMB configuration on the MSDP server with the following precheck script:

```
/usr/opensv/pdde/vpfs/bin/ia_byo_precheck.sh
```

Limitations and considerations

The following limitations and considerations exist for granular restore:

- If adequate space is not available on the target location, the restore operation fails before the copy operation begins.
- The following devices are ignored when snapshots are performed or indexed.
 - Ephemeral storage devices
(For example, Amazon AWS instance store volumes and Microsoft Azure temporary disks.) These devices are also ignored for indexing as well.
 - File systems that are created on the LDM disk
These file systems are ignored for host consistent snapshots.
- Until the old agent (pre-installed) service is not restarted, alternate host restore (GRT and application) of the LVM asset might fail. To support the recovery of LVM assets, you need to restart the older agents.
- Granular restore (GRT) or single file restore (SFR) can be performed with the help of VxMS indexing. VxMS indexing is applicable for all Snapshot Manager-supported file systems. VxMS indexing can be performed for Azure, Azure Stack, AWS, and GCP.
- Host consistent snapshot is supported for the EXT2 file system only if it is mounted as read-only.
- If any unsupported file systems are present on the host, the host can 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.

- During indexing, OS errors can occur while crawling files, directories, or other entries. These errors are ignored and the indexing operation continues. To restore the missing files, you must initiate the granular restore operations on the parent folder.
- When you create or mount a disk from the Windows VM, add the drive letter. This action ensures that the indexing operation can capture the correct drive letter.
- In some cases a mount point is not visible when you browse for files or folders from the recovery point. Consider 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.

Consider the following example. A disk is mounted on `/mnt1/mnt2` where `/mnt1` is any directory on the "/". (The root file system that is on the LVM setup.) `mnt2` is a mount point inside `mnt1`. `mnt2` is not visible in the tree on the left panel. However, you can search and restore files or folders inside the 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 a higher version to a lower version is not recommended.
- A user group cannot restore a drive as a source to an alternate folder as the destination. A user group does not have the writer permission to create a new folder.
- The agentless connection cannot restore the encrypted file by Windows (or EFS) through a granular file-level restore (Restore files and folder option). However, you can restore the file through a volume-level restore and then decrypt the file.
- Files that are stored on a volume that is mounted on a folder (junction point) can be restored only if the underlying disk has the GPT partition layout. If the volume is mounted using a drive letter, then the files can be restored irrespective of the partition layout of the underlying disk.

- If an alternate path that does not exist on the RHEL target host is specified for single file restore, then the new directories that are created are under the security context of the user under which the agent runs. The storage administrators must ensure that the final restore location is accessible to the required user.
- NetBackup does not support indexing of Azure Ultra disks and granular restore.

Limitations for single file restore from a backup copy

- When you restore files and folders from a Linux source host and the target host is Windows, the following points apply:
 - File attributes cannot be restored on a Windows host and only the content of the file is restored.
 - If there is any symlink in the files or folders that are selected for restore, the symlink is not restored.
 - For a restore to the original location, the check for available size is skipped before the copy operation.
- If restoring files or folders when the source host is Linux and the target host is Linux, then the socket and the block files are not restored.
- A restore of files and folders is not supported when they reside on any LDM disks, dynamic disks, or storage spaces.
- If the media server or the PureDisk Deduplication Engine and Veritas Provisioning file system daemon service restarts, the live mount that is retained during a partially successful restore is removed or expires before the retention period expiration date.
- If any media servers are not upgraded to 10.3, then the primary server on version 10.3 is used to connect to NetBackup Snapshot Manager.
- The junction point on Windows after indexing uses the following format:
Volume {4e3f8396-490a-400a-8abf-5579cafd4c0f}
To restore a junction point for single file restore from backup operation, select **Restore everything to a different location** and in the Advanced options enable **Require to restore access control list**.

Operational notes for the Activity monitor

The following behaviors exist for the Activity monitor:

- After a restore job is completed, you cannot expand the directories in the **File List** section of the restore job.

- 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 no longer displays.
- Bytes transferred and estimated bytes are not updated and are shown as 0.

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 **Workloads > 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 jobs for the assets are complete, click the **Recovery points** tab.
- 9 For the preferred recovery point, select **Restore files and folders** from the Action menu.

You can also restore files and folders for **Snapshot** and **Backup** type of copies by clicking **Recover** and then selecting **Restore files and folders**.

- 10 In the Add file step, click **Add**.
- 11 In the **Add files and folders** dialog, 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 In the Recovery target step, perform the following:

Dialog box	Snapshot copy	Backup copy
Restore to	Target VM - Select a VM. A list of all connected VMs having the same operating systems as the original target host is displayed. If you do not select a VM, the files are restored to the original VM.	<ul style="list-style-type: none"> ■ Cloud provider - Select the cloud provider to where single file restore is to be performed. ■ 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. ■ <i>(For Azure and Azure Stack only)</i> Subscription - To restore to an alternate subscription, select one from the drop-down. ■ Target VM - Select a VM. A list of all connected/disconnected and Linux/Windows VMs is displayed for cross-platform restore.
Restore target options	<ul style="list-style-type: none"> ■ Restore everything to original location ■ Restore everything to a different location You must then provide a directory location. You can also enter a UNC path to the location. 	

Restoring files and folders across cloud providers is supported using granular restore from backup copy. Source VM and target VM can be part of different cloud providers to perform granular restore.

Cross-platform restore is supported for the following scenarios:

- NetBackup and Snapshot Manager on one cloud, target host on another cloud.
 - NetBackup and Snapshot Manager on one cloud, another Snapshot Manager and target host on another cloud.
 - NetBackup and Snapshot Manager on one cloud, AIR (Auto Image Replication) restore on another domain.
- 14 If **Restore everything to original location** option is selected, then click **Next** and select the following preferred option in the Recovery options step:

Dialog box	Snapshot copy	Backup copy
Options	<ul style="list-style-type: none">■ Append string to file names In the String field, enter the string that you want to use to append. The string is appended before the last extension of a file.■ Allow overwrite of existing files You must have appropriate permissions.	
Advanced options	N/A	<ul style="list-style-type: none">■ <i>(Applicable only for Windows to Windows restore)</i> Require to restore access control list - Select the checkbox to restore access control list which requires additional operations.■ Target host NAT gateway IP address - Enter network address translation gateway IP address, in case the target VM is behind a network gateway and is not directly accessible. Note: Only private IPs or host names are allowed.

- 15** If **Restore everything to a different location** option is selected, then provide the **Directory for restore** and click **Next**.
- 16** 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 a summary of restored files in the job details.

Note: Permissions on files are assigned based on uid/guid, during restore to non-similar environments (where user/groups do not match). Restored files/folders must have permission for non-intended users/groups on the target host. Hence after a successful restore of the required files, you must modify the access as per requirement.

Note the following:

When restoring hard links for single file restore from a snapshot or a backup (source Linux VM to target Linux VM), ensure the following guidelines:

- When selecting folders and files in the **Add files and folders** dialog, do not select redundant entries. For example, selecting a folder and a file that exists in the same folder, since the folder already has that file.
- Even if redundant entries are selected, ensure that you do not select the **Allow overwrite of existing files** option in the Recovery option step. Selecting this option fails copying the hard link file.

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 **Workloads > 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 to which you want to restore the volume(s).

To restore from a replicated (non-primary) VM, restore to the 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.

Note: If you want to restore the volume to the same virtual machine and location, you must detach the existing volume to free the slot and then try to restore.

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 are performed through the installed agents. However, 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 PV's. If there are duplicate PVs (a warning is displayed on the above command) then run,

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

Otherwise, 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 discovers new VG as <NewVGName>

```
vgimport -a  
  
vgs
```

- 2 Run this 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 authenticated 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 start a subsequent (twice) restore operation on the same VM, an error occurs during a restore operation. This error causes the following issues:

- The tags from the original OS disk are not copied to the newly created restored OS disk.
- User logon 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 in the following article.

learn.microsoft.com/en-us/troubleshoot/azure/virtual-machines/troubleshoot-ssh-connection

Filtering unsupported files and folders

If you try to restore files or folders from a partition or a file system that Snapshot Manager does not support, then you get the following error in the restore job.

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

Workaround:

You can filter any files or folders that Snapshot Manager does not support. In the `bp.conf` file on the primary server, enable the CP DISKMAP check by setting the following flag.

```
CP_DISKMAP_CHECK = true/yes
```

Backup from restore operation is partially successful

Backup from restore operation is partially successful when disk is full on the selected target directory. The following messages are displayed:

```
Dec 29, 2022 2:57:51 PM - Info nbcs (pid=2244) Granular restore(SFR) is completed
Dec 29, 2022 2:57:51 PM - Info nbcs (pid=2244) Summary of SFR Operation - Success
files/folders count: 0 ,
Failed files/folders count: 1 , Warning files/folders
count: 0, Skipped files/folders count: 0
Dec 29, 2022 2:57:51 PM - Info nbcs (pid=2244)
Detailed restore summary report is available on recovery target host at location:
/var/log/flexsnap/restore/granular-restore-09b4d44d
.
```

```
.  
Dec 29, 2022 2:57:51 PM - Warning bprd (pid=1977) Granular Restore from backup  
completed with error.
```

```
Copy the files manually from live access mount:
```

```
ip-10-239-185-241:/mnt/vpfs_shares/vmfiles/8fcc/8fcc132b-a202-49a8-b654-81ff242a718a/livemount
```

```
Dec 29, 2022 2:57:51 PM - end Restore; elapsed time 0:01:51  
the requested operation was partially successful(1)
```

For a backup from restore, if a live mount is created successfully, then even if other errors are reported apart from `ASSET_NOT_FOUND`, it is considered as partial success. If no network devices or a file system are mounted on the target location, or the disk is full, the following messages are displayed in the job details:

```
Jan 02, 2023 12:11:16 AM - Error nbcs (pid=13934)  
187776K space required for file/folder restore while 20K is total available space on  
/disk1
```

In this case, other network devices or the file system must have been mounted on the target path, hence the Snapshot Manager agent considers the free space on the device or file system. As the copy fails with space error, it is logged into the summary report. For example:

```
/var/log/flexsnap/restore/granular-restore-09b4d44d in above Job details log
```

Workaround:

- Check the summary report on the target-host location. For example,

```
/var/log/flexsnap/restore/granular-restore-09b4d44d  
[root@ip-10-239-187-148 granular-restore-09b4d44d]# cat root-error.log  
Dec 29 09:27:44: ERROR - FILE: /disk1/dl380g9-149-vm15_package.zip  
[Error 28] IOError: No space left on device
```

- If the file copy operation failed due to disk space, then create some space and copy the file from live mount.

The live mount path details can be found in the job details as follows:

```
Dec 29, 2022 2:57:51 PM - Warning bprd (pid=1977) Granular Restore from backup completed  
with error.
```

```
Copy the files manually from live access mount:
```

```
ip-10-239-185-241:/mnt/vpfs_shares/vmfiles/8fcc/8fcc132b-a202-49a8-b654-81ff242a718a/livemount
```

Partial recovery is observed when the user selects a disconnected target virtual machine

Partial recovery may be due to the following reasons:

- If the target virtual machine is disconnected (no connectivity through an agent).
- If any failure is seen during the copy of files or folders on the target virtual machine.
- When a Windows virtual machine content is restored on a Linux target virtual machine.

In these partial recovery cases, the created instant access is not deleted and is available for the next 24 hours.

You can configure the instance access retention interval using the **CLOUD_VM_IA_RETENTION_INTERVAL_IN_HOURS** parameter in the `bp.conf` file. (Default value is 24 hours.)

Workaround:

User can perform manual steps to access the instant access share on the target host and then manually copy the required files or folders.

(Copy files over NFS) To restore Linux image content on a Linux host:

- To mount an NFS share on a Linux system, install the NFS client package using the following command:

```
$ sudo yum install nfs-utils
```

- Using the following mount command, mount the Instant access on the target Linux host:

```
# Create a directory say /mnt/restore
```

```
$ mkdir -p /mnt/restore
```

```
# Mount the instant access
```

```
$ mount -t nfs <InstantAccessServer:InstantAccessPath> /mnt/restore
```

- Instant access path can be retrieved from activity manager logs which is in the following format:

```
<InstantAccessServer>:/mnt/vpfs_shares/vmfiles/<id>/<InstantAccessId>/livemount
```

(SMB access) To restore Windows image contents (with ACL) on Windows target host:

- SMB credentials of MSDP storage server of source virtual machine image must be added to the Windows credential manager.

- Use the given live mount to access virtual hard disks by navigating to **Activity Monitor > Job details**.
The virtual hard disks are listed under the folder with **vhd_** as the prefix.
- From the **Action** tab, attach the required virtual hard disk and click on **OK**.
- Select **Assign the following drive letter** option to assign the letter to virtual disk to browse data and click on **OK**.
- Navigate to the assigned drive in the previous step and copy the data manually.

(Live mount) To restore Windows image contents on a Linux target host:

- Linux must have the CIFS package. Obtain the packages using the `# yum install cifs-utils` command.
- Create the mount directory using the `# mkdir <my_mount_dir>` command.
- Use your Samba username and password to mount the exported path as follows:

```
mount -t cifs -o username=<sambauser>
//<InstantAccessServer>/<InstantAccessPath> <my_mount_dir>
```
- Copy the files using the following command:

```
# cp <my_mount_dir>/<file_path> <target_dir_path>
```

Issues with single file restore from the backup of a snapshot

Issue/Error	Description	Workaround
Log path to check	<p>For information related to restore details on the target host, check the following logs:</p> <ul style="list-style-type: none"> ■ <code>path/file:</code> <code>/tmp/flexsnap-agentless-onhost.log</code> ■ <code>/var/log/flexsnap/restore/granular-restore-*</code> 	<p>To resolve the failures or any exceptions that occurred during the single file restore on Snapshot Manager, refer to the following logs on the Snapshot Manager host: <code>/cloudpoint/logs/flexsnap.log</code></p>
Pre-recovery check fails	<p>When restoring files and folders to the disconnected target virtual machine, the pre-recovery check fails with the following error:</p> <pre>Target VM state: Target VM <vm_name> has no agent configured</pre> <p>If recovery is started, the restore operation is partially successful.</p>	<p>Ensure that the target virtual machine is connected to the agent configured for successful restore.</p>

Issue/Error	Description	Workaround
Partial recovery for source Linux VM to target Windows VM (no NFS client)	<p>If you do not install the NFS client on the Windows target computer, a restore of files and folders from a source Linux VM is partially successful. The following error displays:</p> <pre>Error nbcs (pid=42513) Invalid operation for asset: <asset_id> Warning bprd (pid=42045) Granular Restore from backup completed with error. Copy the files manually from live access mount: <livemount_path>. Note that live access mount is available only for 24 hrs.</pre>	Install the NFS client on the Windows target computer before you perform the restore from a Linux VM to a Windows VM.
Restore job fails for deleted target VM	<p>The restore job fails with the following error when restoring files and folders to the target VM which is deleted from the cloud environment:</p> <pre>Error nbcs (pid=44859) Target VM not found, asset_id <asset_id></pre>	Select a different target VM.
Creating instant access fails	<p>If instant access is not enabled on the MSDP storage server, the creation of the instant access fails during the restore job.</p>	<p>Verify if instant access is supported on the MSDP media server. Run the following pre-check script:</p> <pre>/usr/openv/pdde/vpfs/bin/ia_byo_precheck.sh</pre>
Target VM does not have the free drives to attach to the virtual disk	<p>If the number of volumes that contain the selected file(s) is more than the number of free available drives on the target host, the operation fails.</p>	Select a smaller number of volumes for the restore.
Not enough space: "%\driverMapping.json	<p>The media server where MSDP is configured has FIPS enabled.</p>	<p>Disable FIPS on the media server where MSDP is installed. Or, add the domain user Samba credentials to the target VM.</p>

Issue with Azure cloud provider VMs

If any one of the disks in the VM is not initialized, download or restore of the VM's files using instant access fails with the following error:

```
Jan 24, 2023 11:58:47 AM - Error NBWMC (pid=3716) Internal Error:
('failed to find operation system information, please check the source
VM', ('Failed to expose
VMDK', 1006), None)
```



```
Failed to create the instant access mount.  
(4001)
```

`libguestfs` is a third-party tool that instant access uses to retrieve files from a VM backup. If a disk is not initialized, `libguestfs` cannot retrieve the files.

Workaround:

Initialize the disk and back up the VM. Then try again to download or restore the VM files using instant access.

Troubleshooting protection and recovery of cloud assets

This chapter includes the following topics:

- [Troubleshoot cloud workload protection issues](#)
- [Error Code 9855: Error occurred while exporting snapshot for the asset: <asset_name>](#)
- [Backup from snapshot jobs take longer time than expected](#)
- [Backup from snapshot job fails due to connectivity issues when Snapshot Manager is deployed on an Ubuntu host](#)
- [Error disambiguation in NetBackup UI](#)
- [Troubleshoot PaaS workload protection and 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 10.

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

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

Log files for configuration

Use the following logs to troubleshoot cloud configuration issues.

Table 4-1 Log files for configuration

Process	Logs
tpconfig tpconfig command is one way to register Snapshot Manager in NetBackup.	Windows <i>NetBackup install path/volmgr/debug/tpcommand</i> UNIX <i>/usr/opensv/volmgr/debug/tpcommand</i>
nbwebsevice 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/nbwebsevice</i>
nbemm nbemm stores the Snapshot Manager and plug-in information in the 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 4-2 Log files for asset discovery

Process	Logs
ncfnbcs Verifies if the discovery was completed or not.	Windows <i>NetBackup install path/bin/vxlogview -o 366</i> UNIX <i>/usr/openv/netbackup/bin/vxlogview -o 366</i>
Picloud Provides the details of the discovery operation.	Windows <i>NetBackup install path/bin/vxlogview -i 497</i> UNIX <i>/usr/openv/netbackup/bin/vxlogview -i 497</i>
nbwebservice To get details about the asset database workflows that are part of the discovery operation. Note: Refer to the same log files for details of assets that are added to the protection plan.	Windows <i>NetBackup install path/webserver/logs</i> UNIX <i>/usr/openv/wmc/webserver/logs</i> <i>/usr/openv/logs/nbwebservices</i>

Log files for snapshot creation

Use the following logs to troubleshoot snapshot creation issues.

Table 4-3 Log files for snapshot creation

Process	Logs
nbpem nbpem PID for a given job is available in the NetBackup activity monitor.	Windows <i>NetBackup install path/bin/vxlogview -o 116</i> UNIX <i>/usr/openv/netbackup/bin/vxlogview -o 116</i>
nbjm nbjm PID for a given job is available in the NetBackup activity monitor.	Windows <i>NetBackup install path/bin/vxlogview -o 117</i> UNIX <i>/usr/openv/netbackup/bin/vxlogview -o 117</i>

Table 4-3 Log files for snapshot creation (*continued*)

Process	Logs
<p>nbcbs</p> <p>nbcbs PID for a given job is available in the NetBackup activity monitor.</p>	<p>Windows</p> <p><i>NetBackup install path/bin/vxlogview -i 366 -P nbcbs_process_id</i></p> <p>UNIX</p> <p><i>/usr/opensv/netbackup/bin/vxlogview -i 366 -P nbcbs_process_id</i></p> <p>The nbcbs logs are available at the following location:</p> <p>Windows</p> <p><i>NetBackup install path/logs/ncfnbcbs</i></p> <p>UNIX</p> <p><i>/usr/opensv/logs/ncfnbcbs</i></p>
<p>nbrb</p> <p>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 Snapshot Manager.</p>	<p>Windows</p> <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 4-4

Process	Logs
<p>nbwebservice</p> <p>The snapshot restore operation is triggered by NetBackup REST API.</p>	<p>Windows</p> <p><i>NetBackup install path/webserver/logs</i></p> <p>UNIX</p> <p><i>/usr/opensv/wmc/webserver/logs</i></p> <p><i>/usr/opensv/logs/nbwebservices</i></p>
<p>bprd</p> <p>The NetBackup REST API communicates with bprd to initiate restore.</p>	<p>Windows</p> <p><i>NetBackup install path/netbackup/logs</i></p> <p>UNIX</p> <p><i>/usr/opensv/netbackup/logs/bprd</i></p>

Table 4-4 (continued)

Process	Logs
ncfnbcs	Windows
nbcs PID for a 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/opensv/netbackup/bin/vxlogview -i 366 -P nbcs_process_id</i>

Log files for snapshot deletion

Use the following logs to troubleshoot snapshot deletion issues.

Table 4-5 Log files for snapshot deletion

Process	Logs
bpdm	Windows
The snapshot delete or clean-up operation is triggered by bpdm.	<i>NetBackup install path/netbackup/logs</i>
	UNIX
	<i>/usr/opensv/netbackup/logs/bpdm</i>
ncfnbcs	Windows
nbcs PID for a 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/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 the privilege to perform only original location recovery, and you are trying to do alternate location recovery.

Workaround

- While doing the original location restore, do not change any pre-populated fields in the pre-recovery page.

Error Code 9855: Error occurred while exporting snapshot for the asset: <asset_name>

- If you want to perform alternate location recovery, ensure that you have the required privileges.

Error Code 9855: Error occurred while exporting snapshot for the asset: <asset_name>

Explanation:

When you run a backup from snapshot job and the Block Volume Management plug-in is not enabled on the NetBackup Snapshot Manager host, the backup job fails with this error.

Workaround:

Enable the Block Volume Management plug-in for the NetBackup Snapshot Manager host in the **Oracle Cloud Agent** tab of the OCI console.

Backup from snapshot jobs take longer time than expected

Explanation:

Backups from snapshot jobs take longer time than expected when they run with a low transfer rate of around 23 Mbps.

Workaround:

Add the following entry to the `flexsnap.conf` file on the NetBackup Snapshot Manager host.

```
[oci]
```

```
vol_max_vpu_cnt_in_bfs_restore = 120
```

Backup from snapshot job fails due to connectivity issues when Snapshot Manager is deployed on an Ubuntu host

Explanation:

On OCI, when you deploy Snapshot Manager on an Ubuntu host, the default iptable rules may cause issues with network connectivity between the NetBackup services.

These connectivity issues may cause the backup from snapshot, indexing, and restore from backup jobs to fail.

Workaround:

Comment out the iptable rules in the iptable file as shown in the example.

```
Workaround:
    If backup from snapshot needs to be run on Ubuntu deployed
    NBSM (on oracle cloud) then the iptable rules file should look like
    this after commenting out the rules present by default:

    root@nbsm-host:/# cat /etc/iptables/rules.v4
    # CLOUD_IMG: This file was created/modified by the Cloud
    Image build process
    # iptables configuration for Oracle Cloud Infrastructure
    # See the Oracle-Provided Images section in the Oracle
    Cloud Infrastructure
    # documentation for security impact of modifying or
    removing these rule
```

Error disambiguation in NetBackup UI

NetBackup provides you easy troubleshooting options for the errors that you may encounter during various PaaS and Applications processes. For common operations like credential validation, backup, and restore, NetBackup generates a notification that contains the root cause identifier of the error. The notification contains a link to an article that contains the details of the cause and recommended actions.

Troubleshoot PaaS workload protection and recovery issues

Backup fails with error: 3808 Cannot check if the database exists.

You can see the following message in the Activity Monitor:

```
AuthorizationFailed -Message: The client '<clientId>' does not have
authorization to perform action 'Microsoft.Sql/servers/databases/read' over scope
'<resourceId>' or the scope is invalid. If access was recently granted, please
refresh your credentials.
```

Explanation: This error occurs, when the Snapshot Manager and NetBackup are deployed in AKS, and:

- The media server pod node pool is a different node pool from the Snapshot Manager node pool
- Managed Identity is enabled in the Snapshot Manager Virtual Machine Scale set

Workaround: Do any of the following:

- In the media server for backup and restore, enable Managed Identity on the Scale set. Also, assign required permission in the role attached to this managed identity.
- Create a storage unit on the MSDP server, and use only those media servers that have the Managed Identity feature enabled on Scale configuration.

Backup fails when the database or the resource group has a read-only lock applied, and is partially successful when the delete lock is applied.

Explanation: This issue occurs if the Read-only lock or Delete lock attribute is applied to the database or the resource group.

Workaround: Before performing any backup or restore, remove any existing Read-only lock and Delete lock attributes from the database or the resource group.

Status Code 150: Termination requested by administrator

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

Workaround: Manually clean up the database on the provider portal, and temporary staging location at the universal share mount location under a specific directory created by the database name.

Stale status messages in the Activity monitor

Explanation: If the Snapshot Manager 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 Snapshot Manager:

```
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.

Status Code 233: Premature eof encountered

Explanation: Appears if the client name used for the backup exceeds the length of 255 characters.

The `bpdgm` log confirms the same by displaying the following error message:

```
db_error_add_to_file: Length of client is too long. Got 278, but
limit is 255. read_next_image: db_IMAGEreceive() failed: text exceeded
allowed length (225)
```

Note: This is observed when the primary server is RHEL.

Workaround: Rename the database such that the client name fits within the length of 255 characters.

Error: Broken pipe (32), premature end of file encountered EXITING with status 42, network read failed

Or,

Status 174: media manager - system error occurred

Explanation: Occurs during backup if the policy prefix length during protection plan creation is larger than the allowed length. Due to this the file path length of the catalog image exceeds 256 chars and fails with the above error message in activity monitor.

The `bpdgm` log confirms the same by displaying the following error message:

```
<16> db_error_add_to_file: cannot stat(\\?C:\Program Files\Veritas
\NetBackup\db\images \azure-midb-1afb87487dc04ddc8fafa453dccb7ca3+
nbux-qa-bidi-rg+eastus+az-sql-mi-bidinet01+
testdb_bidinet02\1656000000\tmp\catstore\
BACKUPNOW+141a73e7-cdc4-4371-823a-f170447dba2d_
1656349831_FULL.f_imgUserGroupNames0): No such file or directory (2)
<16> ImageReadFilesFile::get_file_size: cannot stat(\\?C:\Program
Files\Veritas\NetBackup\db
\images\azure-midb-1afb87487dc04ddc8fafa453d
ccb7ca3+nbux-qa-bidi-rg+eastus+az-sql-mi-bidinet01+testdb_
bidinet02\1656000000\tmp\catstore\BACKUPNOW+141a73e7-cdc4-4371
-823a-f170447dba2d_1656349831_FULL.f_imgUserGroupNames0): No such
file or directory (2) <16> ImageReadFilesFile::executeQuery: Cannot
copy \\?C:\Program
Files\Veritas\NetBackup\db\images\azure-midb-1afb87487dc04ddc8fafa453dccb7
ca3+nbux-qa-bidi-rg+eastus+az-sql-mi-bidinet01+testdb_bidinet02\1
```

```
656000000\tmp\catstore\BACKUPNOW+141a73e7-cdc4-4371-823a-f170447d
ba2d_1656349831_FULL.f_imgUserGroupNames0
```

Note: This is observed when the primary server is Windows.

Workaround: Use a policy prefix name in protection plan with length less than 10 characters, so that the total length of the catalog path is less than 256 characters.

Status Code 3801: Cannot complete the requested operation.

Explanation: NetBackup is not able to successfully carry out the requested operation.

Recommended action: Refer to the activity monitor details for the possible reasons for failure.

Status Code 3817: Cannot complete the pre-backup operation

Explanation: The error message seen in dbagentsutil logs as, pg_dump: error: query failed: ERROR: permission denied for table test; pg_dump: error: query was: LOCK TABLE public.test IN ACCESS SHARE MODE; Invoked operation: PRE_BACKUP failed

Occurs when you try to back up a database that has multiple tables with different roles. If tables have at least one different owner, other than the database owner, and it is not a member of the database owner role, then the backup may fail.

Recommended action: You must have a role that has access to all tables inside the database that you want to backup or restore.

For example, say that we wanted to back up the `School` database which has two tables.

- `student`, owner is `postgres`
- `teacher`, owner is `schooladmin`

Create a new role. Say, `NBUbackupadmin`

Run the following command to create the role:

```
postgres=> CREATE USER NBUbackupadmin WITH PASSWORD '*****';
CREATE ROLE
```

To make this new role a member of `postgres` and `schooladmin` role, run:

```
postgres=> GRANT postgres TO NBUbackupadmin;
GRANT ROLE
```

```
postgres=> GRANT schooladmin TO NBUBackupadmin;

GRANT ROLE
```

Note: You must have a role who is either owner or member of the owner of the table, for all tables inside the database.

Backup fails with Status 40 (Network connection broken)

Explanation: Backups fail due to loss of connectivity to the media server.

Recommended action: You can restart the backup job if the policy has checkpoints enabled. Once the network issue is resolved, select the incomplete backup job in the web UI and click **Resume**. The job resumes from the point it was stopped. If the checkpoint is not enabled in the policy, the job shows up as a failed job in the web UI.

Backup job fails with the error: "Failed to backup database"

Explanation: The Job details contain additional details: ManagedIdentityCredential authentication unavailable. The requested identity is not assigned to this resource. The allocated media server does not have any Managed Identity attached to it.

Recommended action: If you use system or user-managed identity for the PaaS Azure SQL and Managed Instances, apply the same set of permissions/rules to the media server(s) and the snapshot manager. If you use user-managed identity, attach the same user-managed identity to the media server(s) and the snapshot manager.

Error code 3842 - The requested backup type for the corresponding PaaS asset is unsupported.

Differential incremental backup is supported for only for Azure SQL server and Azure SQL Managed Instance. When you select an unsupported backup type, this error appears.

Error code 3843 or 3844 - Failed to enable or disable CDC.

Appears when you do not have permissions to enable or disable the CDC.

Workaround: Give NetBackup the necessary permissions to enable or disable the CDC in your Azure environment.

Note: Do not enable CDC manually. Provide the permissions to NetBackup to enable or disable the CDC.

Error: Client restore EXIT STATUS 5: the restore failed to recover the requested files Cloud policy restore error (2824)

Error: ERR - Failed to restore database [<db_name>] with name [<db_name>]. ERR - Failed to open file". Errno = 12: client restore EXIT STATUS 5: the restore failed to recover the requested files

Explanation: Occurs during restore if the backup image is generated on 10.2 media and restore goes to an older (< 10.2) media server.

Workaround: Change the restore media to 10.2 and remove the older media from storage.

AWS DynamoDB table does not have auto scaling enabled, after restoring from a backup image that has auto scaling enabled

Explanation: Currently the AWS API response does not show if a table has auto scaling enabled. So, during backup, this metadata is not captured in NetBackup, and as a result the restored table does not have auto scaling enabled.

Workaround: Enable the auto-scaling property of the restored DynamoDB table in the AWS portal manually.

CDC-enabled Azure SQL MI incremental backups: Dropping a CDC-enabled database leads to full backup without schema changes, instead of incremental.

Explanation: Azure SQL MI maintains CDC-enabled database details in the table `cdc_jobs` inside the `msdb` schema. When the database is dropped, its `cdc_jobs` entry should be deleted. Sometimes this entry does not get deleted from the `cdc_jobs` table. So, when a new database is created with the same `db_id` which already exists in the `cdc_jobs` table, the issue occurs.

Workaround: When you drop a database, check the entry of the dropped database in the `cdc_jobs` table of the `msdb` schema. If the entry is present there, delete it manually.

AWS RDS: Error while fetching details of db instance: An error occurred (SignatureDoesNotMatch) when calling the DescribeDBInstances operation: Signature expired.

Explanation: Failure of the RDS boto3 APIs shows this error. NetBackup shows this error for the DescribeDBInstances operation.

Workaround: Synchronize the date and time of the media server with the actual network date and time.

Also, verify if you are using the correct provider credentials.

Troubleshooting Amazon Redshift issues

Restore fails for Amazon Redshift, if the query string is larger than 100 KB

Explanation:

This is a known limitation of AWS. The maximum query statement size is 100 KB. See the AWS documentation for the details:

<https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html>

After a successful Redshift database restore, if the number of stored procedures, views, and functions are not the same as the source database.

Workaround:

Do the following:

- 1 Mount the Instant Access (IA) path using the following API:

```
netbackup/recovery/workloads/cloud/paas/instant-access-mounts
```

- 2 Navigate to the mount path in the media server.
- 3 Ensure that the mount path directory hierarchy is as follows:

```
ClusterDirectory/DatabaseDirectory/DatabaseDirectory/SchemaDirectory/TableDirectory
```

- 4 In the `SchemaDirectory`, locate the files `StoredProcedures.json`, `Views.json`, and `Functions.json`. Each file contains one or more SQL statements which you can run in Amazon Redshift Query Editor-2.

Manually run these SQL statements.

botocore.exceptions.ClientError: An error occurred (InvalidSignatureException) when calling the ListDatabases operation

Explanation:

If the system time where you run the AWS Redshift APIs is not correct, you get this error. This message appears in the logs:

```
Signature expired: 20230226T181919Z is now earlier than
20230226T181921Z (20230226T182421Z - 5 min.)"
```

Workaround:

Run the `ntpdate` command to fix the system time.

Backup or restore jobs fail with the "NoCredentialsError: Unable to locate credentials" error.

Explanation:

This error appears when the region is not specified. You can see the following error in the `dbagentsutil` logs. You can find the `dbagentsutil` logs at the following location:

```
/usr/opensv/netbackup/logs/
```

Workaround:

Do the following:

- 1 Download AWS CLI on the media server where the `dbagent` is running.
- 2 Run the command:


```
aws configure
```
- 3 Enter the region name for EC2 when prompted. Do not specify the values for the other parameters.

Backup and restore stuck for Redshift databases

Explanation:

This error appears when the NetBackup Snapshot Manager that runs the discovery does not have access to the Redshift cluster. You can see the following error in the `flexsnap` logs:

```
Connect timeout on endpoint URL:
"https://redshift.us-east-2.amazonaws.com/
```

Workaround:

Without access permission, the Snapshot Manager requires the inbound rules to be configured for the snapshot manager in the security group of the 'VPC endpoint of the Redshift service'.

On the AWS portal, select a cluster. Click Properties > click Network and security settings > click the virtual private cloud object > click Endpoints. Search for "redshift-endpoint" in the search field > click the VPC endpoint ID > click the Security Groups tab. Click the Security Group ID > click Edit Inbound rules, and add the following for media servers.

Type : HTTPS

Protocol : TCP

Port range : 443

Source : 10.177.77.210/32

* Here, the source refers to the media server instance.

Run discovery from NetBackup web UI again.