

NetBackup™ Web UI Cloud Administrator's Guide

Release 11.2

Last updated: 2026-05-28

Legal Notice

Copyright © 2026 Cohesity, Inc. All rights reserved.

© 2026 Cohesity, Inc. All Rights Reserved. Cohesity, the Cohesity Logo and other Cohesity Marks are trademarks of Cohesity, Inc. in the US and/or internationally. The information supplied herein is the confidential and proprietary information of Cohesity and may only be used (a) by the intended recipients and (b) in conjunction with validly licensed Cohesity software and services. Find the terms of Cohesity licenses at www.cohesity.com/agreements.

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. COHESITY 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.

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

- Go to [Cohesity Support](#), to search in our knowledge base; or contact us by phone - United States and Canada: 1-855-9CO-HESI (926-4374), option 2.
- Log in to the [Cohesity Support Portal](#) to create a new case.
- Click the (?) icon on the Cohesity UI and select Support Portal.

Support/Service Assistance

First, contact the Service Provider that you have contracted for service and support. If you work directly with Cohesity and have a product warranty/entitlement, repair pricing, or technical support-related question, see your options below:

- To find solutions to your product issues or for suggestions or best practices, visit the [Cohesity Knowledge Base](#).
- Log in to the [Cohesity Support Portal](#) to create a new case.
- To monitor your open cases, log in to the portal and click the **Cases** tab on the home page. This page should have all the case statuses and updates. You can also view individual case status.

Cohesity Software Running on Partner Hardware

For Cohesity software running on qualified third-party hardware, the following support workflow applies:

1. The customer may contact Cohesity Support first if the issue cannot be determined as a hardware issue.

Note: Cohesity cannot process hardware replacement requests for partner hardware.

2. Cohesity Support triages the issue. If it is a software issue, Cohesity Support continues to work on it.
3. If it is a hardware/firmware issue or is suspected to be a hardware/firmware issue, Cohesity provides information about the issue to the customer and requests that the customer open a support ticket with the appropriate partner.
4. If needed, Cohesity Support can join a three-way call with the partner and the customer.
5. The customer informs Cohesity Support on the progress of the partner's case.

Contents

| | | |
|------------------|--|----------|
| Chapter 1 | Managing and protecting cloud assets | 9 |
| | About protecting cloud assets | 10 |
| | Limitations and considerations | 12 |
| | AWS and Azure government cloud support | 13 |
| | Configure Snapshot Manager in NetBackup | 14 |
| | Add a Snapshot Manager | 15 |
| | Add a cloud provider for a Snapshot Manager | 15 |
| | Associate media servers with a Snapshot Manager | 22 |
| | Discover assets on Snapshot Manager | 22 |
| | Enable or disable a Snapshot Manager | 24 |
| | (Optional) Add the Snapshot Manager extension | 24 |
| | Managing intelligent groups for cloud assets | 24 |
| | Considerations for cloud intelligent groups | 25 |
| | Create an intelligent group for cloud assets | 26 |
| | Delete an intelligent group for cloud assets | 30 |
| | Protecting cloud assets or intelligent groups for cloud assets | 30 |
| | Customize or edit protection for cloud assets or intelligent groups | 32 |
| | Remove protection from cloud assets or intelligent groups | 33 |
| | About storage lifecycle policies | 33 |
| | Adding an SLP | 34 |
| | SLP configurations for PaaS and IaaS policies | 37 |
| | Managing policies for cloud assets | 39 |
| | Limitations and considerations | 39 |
| | Planning for policies | 39 |
| | Creating policies for cloud assets | 41 |
| | Setting up attributes for PaaS assets | 42 |
| | Setting up attributes for IaaS assets | 44 |
| | Creating schedules | 46 |
| | About backup frequency | 48 |
| | About assigning retention periods | 49 |
| | Configuring the Start window | 52 |
| | Configuring the include dates | 54 |
| | Configuring the exclude dates | 55 |
| | Configuring the cloud assets for PaaS | 56 |
| | Configuring the cloud assets for IaaS | 58 |

| | |
|---|-----------|
| Configuring backup options for IaaS | 59 |
| Managing cloud policies | 60 |
| Scan for malware | 62 |
| Scanning backup images | 62 |
| Assets by workload type | 64 |
| Protecting Microsoft Azure resources using resource groups | 65 |
| Before you begin | 66 |
| Limitations and considerations | 67 |
| About resource group configurations and outcome | 67 |
| Troubleshoot resource group permissions | 70 |
| NetBackup Accelerator for cloud workloads | 71 |
| How the NetBackup Accelerator works with virtual machines | 71 |
| Considerations for using Accelerator with OCI | 72 |
| Accelerator forced rescan for virtual machines (schedule attribute) | 73 |
| Accelerator backups and the NetBackup catalog | 73 |
| Accelerator messages in the backup job details log | 74 |
| Configuring backup schedules for cloud workloads using protection plan | 74 |
| Backup options for cloud workloads | 77 |
| AWS Snapshot replication | 79 |
| Configure AWS snapshot replication | 80 |
| Using AWS snapshot replication | 82 |
| Support matrix for account replication | 85 |
| Protect applications in-cloud with application-consistent snapshots | 87 |
| Protecting AWS or Azure VMs for recovering to VMware | 88 |
| Cloud asset cleanup | 89 |
| Cloud asset filtering | 90 |
| Chapter 2 Protecting PaaS assets | 93 |
| Protecting PaaS assets | 94 |
| Steps to protect PaaS assets | 94 |
| Protecting PaaS assets in AWS | 94 |
| Protecting PaaS assets in Azure | 97 |
| Protecting PaaS assets in GCP | 100 |
| Prerequisites for protecting PaaS assets | 103 |
| Enabling binary logging for MySQL and MariaDB databases | 106 |
| Enabling backup and restore in Kubernetes | 106 |
| Prerequisites for protecting Amazon RDS SQL Server database assets | 107 |
| Protecting RDS Custom instances | 108 |

| | |
|---|-----|
| Protecting RDS Custom for SQL Server assets | 108 |
| Consideration for protecting RDS Custom for SQL Server assets | 108 |
| Protecting RDS Custom for Oracle assets | 109 |
| Consideration for protecting RDS Custom for Oracle assets | 109 |
| Protecting Azure Managed Instance databases | 110 |
| Prerequisites for protecting Azure Managed Instance databases | 110 |
| Permissions required for protecting Azure Managed Instance databases | 111 |
| Prerequisites for database level discovery | 111 |
| Limitation and considerations | 112 |
| For all databases | 112 |
| For PostgreSQL | 113 |
| For incremental backups for Azure PostgreSQL | 114 |
| For Amazon RDS PostgreSQL and Amazon Aurora PostgreSQL | 114 |
| For Amazon DynamoDB | 115 |
| For Amazon DocumentDB | 115 |
| For Amazon Neptune | 115 |
| For Amazon RDS SQL | 116 |
| For Azure, Amazon RDS, and Aurora MySQL | 116 |
| For cumulative incremental backups using Amazon RDS SQL | 117 |
| For incremental backups using Amazon DynamoDB | 118 |
| For incremental backups using Azure MySQL server | 118 |
| For incremental backups using the GCP SQL Server | 119 |
| For Incremental backups using GCP PostgreSQL | 120 |
| For Incremental backups with Amazon RDS PostgreSQL and Amazon Aurora PostgreSQL | 121 |
| For Azure SQL and SQL Managed Instance | 122 |
| For Azure SQL Server and SQL Managed Instance incremental backup | 124 |
| For Azure Cosmos DB for MongoDB | 125 |
| For Azure Cosmos DB for NoSQL | 125 |
| For Amazon RDS for Oracle | 126 |
| For Amazon Redshift databases | 127 |
| For Amazon Redshift clusters | 127 |
| For GCP SQL Server | 128 |
| For GCP BigQuery | 128 |
| Installing the native client utilities | 129 |
| Installing the MySQL client utility | 129 |
| Installing the SqlPackage client utility | 131 |

| | |
|--|-----|
| Installing PostgreSQL client utility | 133 |
| Installing MongoDB client utility | 135 |
| Configuring storage for different deployments | 136 |
| For MSDP cloud deployments | 136 |
| For Kubernetes deployments | 136 |
| For VM-based BYO deployments | 136 |
| Configuring the storage server for instant access | 137 |
| About incremental backup for PaaS workloads | 138 |
| Configuring incremental backups for Azure MySQL server | 138 |
| About archive redo log backup for PaaS workloads | 140 |
| About Auto Image Replication for PaaS workloads | 140 |
| Discovering PaaS assets | 141 |
| Viewing PaaS assets | 142 |
| Managing PaaS credentials | 142 |
| View the credential name that is applied to a database | 143 |
| Add credentials to a database | 143 |
| Creating an IAM database username | 145 |
| Creating a system or user-managed identity username | 146 |
| Configuring permissions for the database user | 149 |
| Add protection to PaaS assets | 150 |
| Perform backup now | 150 |

| | | |
|------------------|--|------------|
| Chapter 3 | Recovering cloud assets | 152 |
| | Recovering cloud assets | 152 |
| | About the pre-recovery check for VMs | 153 |
| | Supported parameters for restoring cloud assets | 153 |
| | Restoring to a different cloud provider | 154 |
| | Recovering virtual machines | 155 |
| | Recovering applications and volumes to their original location | 159 |
| | Recovering applications and volumes to an alternate location | 159 |
| | Additional steps required after restoring an AWS RDS database instance | 160 |
| | Recovery scenarios for GCP VMs with read-only volumes | 161 |
| | (GCP only) Restoring virtual machines and volumes using the autoDelete disk support | 162 |
| | Perform rollback recovery of cloud assets | 163 |
| | Restore to a different cloud provider | 163 |
| | Prepare the VMs for back up | 165 |
| | Post-restore configurations | 170 |
| | Recovering AWS or Azure VMs to VMware | 173 |

| | | |
|------------------|--|------------|
| | Post-recovery considerations for cloud VMs recovered to VMware | 175 |
| | Steps to recover images from cloud VMs to VMware | 175 |
| | Recovering PaaS assets | 178 |
| | Recovering non-RDS PaaS assets | 178 |
| | Recovering Redshift clusters | 179 |
| | Recovering AWS DocumentDB and Neptune assets | 180 |
| | Recovering RDS-based PaaS asset | 181 |
| | Recovering Azure-protected assets | 182 |
| | Recovering duplicate images from AdvancedDisk | 184 |
| Chapter 4 | Performing granular restore | 186 |
| | About granular restore | 186 |
| | Supported environment list | 187 |
| | List of supported file systems | 189 |
| | Before you begin | 189 |
| | Limitations and considerations | 191 |
| | Restoring files and folders from cloud virtual machines | 195 |
| | Restoring volumes on cloud virtual machines | 199 |
| | Performing steps after volume restore containing LVM | 200 |
| | Troubleshooting | 202 |
| Chapter 5 | Troubleshooting protection and recovery of cloud assets | 209 |
| | Troubleshoot cloud workload protection issues | 209 |
| | Error Code 9855: Error occurred while exporting snapshot for the asset: <asset_name> | 214 |
| | VMs and other OCI assets with CMK-encrypted disks are marked as deleted in NetBackup UI. | 214 |
| | Backup from snapshot jobs take longer time than expected | 214 |
| | Backup from snapshot job fails due to connectivity issues when Snapshot Manager is deployed on an Ubuntu host | 215 |
| | Error disambiguation in NetBackup UI | 215 |
| | Status Code 150: Termination requested by administrator | 216 |
| | Troubleshoot PaaS workload protection and recovery issues | 216 |
| | Troubleshooting Amazon Redshift issues | 223 |
| | Troubleshooting Azure Postgres issues | 226 |
| | Troubleshooting Amazon RDS Custom for SQL issues | 226 |

Managing and protecting cloud assets

This chapter includes the following topics:

- [About protecting cloud assets](#)
- [Limitations and considerations](#)
- [AWS and Azure government cloud support](#)
- [Configure Snapshot Manager in NetBackup](#)
- [Managing intelligent groups for cloud assets](#)
- [Protecting cloud assets or intelligent groups for cloud assets](#)
- [About storage lifecycle policies](#)
- [Managing policies for cloud assets](#)
- [Scan for malware](#)
- [Protecting Microsoft Azure resources using resource groups](#)
- [NetBackup Accelerator for cloud workloads](#)
- [Configuring backup schedules for cloud workloads using protection plan](#)
- [Backup options for cloud workloads](#)
- [AWS Snapshot replication](#)
- [Protect applications in-cloud with application-consistent snapshots](#)
- [Protecting AWS or Azure VMs for recovering to VMware](#)

- [Cloud asset cleanup](#)
- [Cloud asset filtering](#)

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 |
|--|--|
| <p>Before you begin, ensure that you have the appropriate permissions.</p> | <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> |
| <p>Deploy Snapshot Manager</p> <p>Configure the Snapshot Manager</p> | <p>Install Snapshot Manager in your environment.</p> <p>See "Add a Snapshot Manager" on page 15.</p> <p>Review Snapshot Manager and NetBackup limitations.</p> <p>See "Limitations and considerations" on page 12.</p> <hr/> <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 15.</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 22.</p> |
| Create a protection plan or policy | <p>Create a protection plan or policy. 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 80.</p> <p>See “Managing policies for cloud assets” on page 39.</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 87.</p> |
| Manage policies for cloud assets | <p>Using the NetBackup Web UI, you can create policies that support cloud workloads types (IaaS and PaaS). The policies are applied on the workloads to protect the data that resides on the clients.</p> <p>See “Managing policies for cloud assets” on page 39.</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 152. See “Recovering cloud assets” on page 152. See “Perform rollback recovery of cloud assets” on page 163. ■ 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. |
| Support for malware scan before recovery | You can trigger malware scan of the selected files/folders for recovery as part of recovery flow from Web UI and decide the recovery actions based on malware scan results. |
| Troubleshooting | See “Troubleshoot cloud workload protection issues” on page 209. |

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, Cohesity 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: Cohesity 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.

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 do not support government cloud.

For information about configuring AWS and Microsoft Azure, See [“Add a cloud provider for a Snapshot Manager”](#) on page 15.

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 Hub, GCP and from version 10.4 OCI cloud is also supported.

Consider the following important points:

- You can associate multiple Snapshot Manager servers to a NetBackup master server. But Snapshot Manager server cannot have multiple NetBackup master servers.
- You can now manage Snapshot Manager and control the discovery of assets from the NetBackup web UI, REST APIs, and CLIs 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 15. |
| 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 15. |
| Discover assets on Snapshot Manager | You can discover assets on the Snapshot Manager. See “Discover assets on Snapshot Manager” on page 22. |
| 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 22. |

Add a Snapshot Manager

You can add a Snapshot Manager using the NetBackup web UI.

Starting with NetBackup 11.1, a cloud administrator can use a Snapshot Manager configured by a NAS administrator to manage storage arrays across different cloud providers, such as Azure Files, Azure NetApp Files, and FSx for AWS.

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**.

The added Snapshot Manager is also visible in the **Storage > Snapshot Manager** tab.

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. Note: The user-managed identity must have a role assigned. | |
| | <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 or US Gov Virginia. | |
| | Resource Group prefix | The prefix used to store the snapshots created for the assets in a different resource group other than the one in which the assets exist. For example, if an asset exists in NetBackup Snapshot Manager and prefix for resource group is snap , then snapshots of assets in NetBackup Snapshot Manager resource group would be stored in snapNetBackup Snapshot Manager resource group. | |
| Protect assets even if prefixed Resource Groups are not found | On selecting this check box, NetBackup Snapshot Manager would not fail the snapshot operation if resource group does not exist. It tries to store the snapshot in the original resource group. Note: The prefixed resource group region must be same as the original resource group region. | | |

| Cloud provider | Parameter | Description |
|---|--|--|
| Microsoft Azure Stack Hub | <i>Using AAD:</i> | |
| | Azure Stack Hub Resource Manager endpoint URL | The endpoint URL in the following format allows Snapshot Manager to connect with your Azure resources. <code>https://management.<location>.<FQDN></code> |
| | Tenant ID | The ID of the AAD directory in which you created the application. |
| | Client ID | The application ID. |
| | Secret Key | The secret key of the application. |
| | Authentication Resource URL (optional) | The URL where the authentication token is sent to. |
| | <i>Using ADFS:</i> | |
| | Azure Stack Hub Resource Manager endpoint URL | The endpoint URL in the following format that allows Snapshot Manager to connect with your Azure resources. <code>https://management.<location>.<FQDN></code> |
| | Tenant ID | The ID of the AAD directory in which you created the application. |
| | Client ID | The application ID. |
| Secret Key | The secret key of the application. | |
| Authentication Resource URL (optional) | The URL where the authentication token is sent to. | |

| Cloud provider | Parameter | Description |
|-----------------------|--|---|
| Amazon AWS | Access Key | <p>The access key ID, when specified with the secret access key, authorizes Snapshot Manager to interact with the AWS APIs.</p> <p>Note: For more information on how to create an IAM role, see the AWS documentation.</p> |
| | 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 | <p>One or more AWS regions in which to discover cloud assets.</p> <p>Note: If you configure a government cloud, select us-gov-east-1 or us-gov-west-1.</p> |
| 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 | <p>The private key. Listed as <code>private_key</code> in the JSON file.</p> <p>Note: You must enter this key without quotes. Do not enter any spaces or return characters at the beginning or end of the key.</p> |
| 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 | <p>NetBackup Snapshot Manager must be a part of a dynamic group and that dynamic group must have enough permissions.</p> <p>Note: If the Snapshot Manager is configured with IAM Configuration, the other fields, except Regions, are not available.</p> |

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

7 Click **Save**.

The assets on the cloud providers are automatically discovered.

Adding a new region

You can add new regions to the Snapshot Manager by editing the configuration.

To add a new region:

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

2 Click the **Snapshot manager** tab. Click the tab of the provider for which you want to add one or more regions. Click the ellipsis menu (three dots) in the line of the plug-in to which you want to add the new region.

Or,

Click the **Providers** tab, click **Configurations** for the provider for to which you want to add the region. Click the ellipsis menu (three dots) in the line of the plug-in to which you want to add the new region.

- 3 Add one or more new regions from the **Regions** list.
- 4 Click **Save**, and wait till the **Discovery status** column shows **Success**, in the plug-in properties page.
- 5 Click the ellipsis menu (three dots) in the line of the plug-in, and click **Discover**. Wait for the discovery to complete.

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 15.

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

Note: If you change the IAM role for the NetBackup Snapshot Manager host after the AWS CSP configuration, you need to edit the CSP configuration, and save it once with the same configuration.

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 policy for OCI configuration

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

See [“Add a cloud provider for a Snapshot Manager”](#) on page 15.

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

OCI supports the Source account implementation of IAM policies. Snapshot Manager supports IAM policy configuration for OCI in the same tenancy where the Snapshot manager is deployed. So, OCI cloud is aware of the OCI Tenancy ID, you only need to 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.
- Intelligent groups are not supported for AWS DocumentDB and AWS Neptune workloads.
- 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.

- Intelligent groups are not supported on Redshift clusters. However, intelligent groups for Redshift database assets are supported.
- For Azure MySQL assets, you cannot create intelligent groups with a mix of database and server assets. An intelligent group can either contain a group of databases or servers. When creating an intelligent group for Azure MySQL, you must specify the entityType filter as either server or database.
- 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.

Considerations for Application intelligent groups

While creating an intelligent group for applications, only RDS assets are supported in AWS.

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.

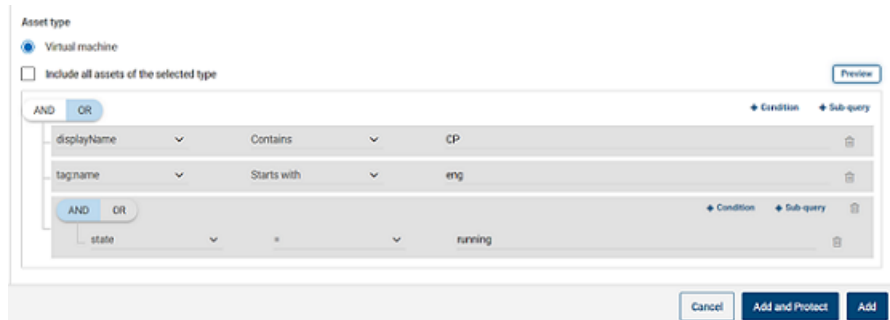
Note: If the region is not specified, then the cloud intelligent group filter applies to the assets from all the discovered regions.

- 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**.

- 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 28.

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



This example uses **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/shape | 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. |
| compartmentId | The asset's compartment OCID. 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 31. 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**.

About storage lifecycle policies

A storage lifecycle policy (SLP) is a storage plan for a set of backups. You can configure SLPs from the NetBackup UI. To view the existing SLPs or create a new one, on the left navigation pane, click **Storage > Storage Lifecycle Policies**.

An SLP contains instructions in the form of storage operations, to be applied to the data that is backed up by a backup policy. Operations are added to the SLP that determine how the data is stored, copied, replicated, and retained. NetBackup retries the copies as necessary to ensure that all copies are created.

SLPs offer the opportunity for users to assign a classification to the data at the policy level. A data classification represents a set of backup requirements, which

makes it easier to configure backups for data with different requirements. For example, email data and financial data.

SLPs can be set up to provide staged backup behavior. They simplify data management by applying a prescribed behavior to all the backup images that are included in the SLP. This process allows the NetBackup administrator to leverage the advantages of different backups in the near term or long term.

This section briefly describes the SLPs, for more details see *NetBackup™ Administrator's Guide, Volume I*.

For SLP best practices, see the Knowledge Article:
<https://support.cohesity.com/s/article/article-100009913>.

Adding an SLP

The operations in an SLP are the backup instructions for the data. Use the following procedure to create an SLP that contains multiple storage operations.

This section briefly describes SLP creation, for more details see *NetBackup™ Administrator's Guide, Volume I*.

For SLP best practices, see the Knowledge Article:
<https://support.cohesity.com/s/article/article-100009913>.

To create an SLP

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Storage lifecycle policies**.
- 3 Click **Add** to create a new SLP.
- 4 On the **Storage lifecycle policy** pane, provide the following details:
 - **Storage lifecycle policy name:** The name cannot be modified after the SLP is created.
 - **Data classification:** Defines the level or classification of data that the SLP is allowed to process. The dropdown menu contains all of the defined classifications as well as the **Any** classification, which is unique to SLPs. The **Any** selection indicates to the SLP that it should preserve all images that are submitted, regardless of their data classification.
 - **Priority for secondary operations:** The priority that jobs from secondary operations have in relationship to all other jobs. The priority applies to the jobs that result from all operations except for Backup and Snapshot operations. Range: 0 (default) to 99999 (highest priority).

For example, you may want to set the **Priority for secondary operations** for a policy with a gold data classification higher than for a policy with a silver data classification.

- 5 Add one or more operations to the SLP. The operations are the instructions for the SLP to follow and apply to the data that is specified in the backup policy. Click **Add** to add operations to the SLP. Provide the following on the **New operation** pane. Select an **Operation** type.

Source storage > Operation

Backup

Snapshot

Import

Destination storage attributes > Destination storage

- Snapshot

- No storage unit

- Snapshot

Destination storage attributes > Volume pool

NetBackup

Note: For Snapshot and Import operation, this option is disabled.

Retention > Retention type

- The **Fixed** retention indicates that the data on the storage is retained for the specified length of time, after which the backups or snapshots are expired. Expires immediately, 1 week, 2 weeks, 3 weeks and many more. An image copy with a fixed retention is eligible for expiration when all of the following criteria are met:
 - The Fixed retention period for the copy has expired.
 - All child copies have been created.
 - All child copies that are mirror copies are eligible for expiration.
- The **Expire after copy** retention indicates that after all direct (child) copies of an image are successfully duplicated to other storage, the data on this storage is expired. The last operation in the SLP cannot use the **Expire after copy** retention type because no subsequent copy is configured. Therefore, an operation with this retention type must have a child.
- The **Capacity managed** operation means that NetBackup automatically manages the space on the storage, based on the High water mark setting for each volume. The High water mark and Low water mark settings on the disk storage unit or disk pool determine how the space is managed.

To add a child operation, select an operation and then click **Add child**. Select an **Operation** type. For a child operation, the SLP displays only those operations that are valid based on the parent operation that you selected.

- 6 The **Window** tab displays for the available operation types. Use them to specify when the secondary operation runs, create a window for the operation.
- 7 Optionally, select **Postpone creation of this copy until the source copy is about to expire**.

- 8 Under **Advanced**, specify if NetBackup should process active images after the window closes.
- 9 Under **Duplication**, you can allow an alternate read server to read a backup image originally written by a different media server.

To understand the different SLP configurations for various snapshot and backup operations:

See “ [SLP configurations for PaaS and IaaS policies](#)” on page 37.

SLP configurations for PaaS and IaaS policies

For the cloud policy type, it is recommended to create an operation hierarchy for SLP. To help understand different SLP configurations for various snapshot and backup operations along with the combination of backup options, refer to the table. It also provides a distinction between the use cases that are run using the protection plan.

Table 1-5 Protection Plan vs Policy SLP for IaaS policy type

| Protection plan | Protection plan backup option | Equivalent SLP operation for policy | Policy backup options |
|--|--|---|--|
| A.I.R (Replication) | N/A | <ul style="list-style-type: none"> ■ Snapshot <ul style="list-style-type: none"> ■ Backup From Snapshot ■ Replication | N/A |
| Backup from snapshot + Granular recovery | Enable granular recovery for files and folder. | <ul style="list-style-type: none"> ■ Snapshot <ul style="list-style-type: none"> ■ Backup From Snapshot | Enable granular recovery for files and folder. |
| Backup from snapshot + Initiate backup when snapshot is about to expire. | Enable granular recovery for files and folder. | <ul style="list-style-type: none"> ■ Snapshot <ul style="list-style-type: none"> ■ Cloud Snapshot Index ■ Backup From Snapshot (Select "Postpone creation of this copy until source copy is about to expire") | Enable granular recovery for files and folder. |

Table 1-5 Protection Plan vs Policy SLP for IaaS policy type (*continued*)

| Protection plan | Protection plan backup option | Equivalent SLP operation for policy | Policy backup options |
|---|---|---|---|
| Duplicate copy | N/A | <ul style="list-style-type: none"> ■ Snapshot <ul style="list-style-type: none"> ■ Backup From Snapshot ■ Duplication | N/A |
| Initiate a backup when snapshot is about to expire. | N/A | <ul style="list-style-type: none"> ■ Snapshot <ul style="list-style-type: none"> ■ Backup From Snapshot (Select "Postpone creation of this copy until source copy is about to expire") | N/A |
| Keep backup only | N/A | <ul style="list-style-type: none"> ■ Snapshot (Select "Expire after copy" retention) <ul style="list-style-type: none"> ■ Backup From Snapshot | N/A |
| Keep snapshot along with backup | N/A | <ul style="list-style-type: none"> ■ Snapshot <ul style="list-style-type: none"> ■ Backup From Snapshot | N/A |
| Keep snapshot only | N/A | Snapshot | N/A |
| Snapshot only + Granular recovery | Enable granular recovery for files and folder | <ul style="list-style-type: none"> ■ Snapshot <ul style="list-style-type: none"> ■ Cloud Snapshot Index | Enable granular recovery for files and folder |

Table 1-6 Protection Plan vs Policy SLP for PaaS policy type

| Protection Plan | Equivalent SLP operation for policy |
|----------------------|---|
| Backup | Backup as primary operation |
| A.I.R. (Replication) | <ul style="list-style-type: none"> ■ Backup <ul style="list-style-type: none"> ■ Replication |

Table 1-6 Protection Plan vs Policy SLP for PaaS policy type (*continued*)

| Protection Plan | Equivalent SLP operation for policy |
|-----------------|---|
| Duplicate | <ul style="list-style-type: none"> ■ Backup ■ Duplication |

Managing policies for cloud assets

Backup policies provide the instructions that NetBackup follows to back up the workloads. Using the NetBackup web UI, you can create policies to support cloud workloads types: IaaS and PaaS. The policies are applied on the workloads to protect the data that resides on the clients.

Using the policy utility that resides on the clients, you can configure multiple types of policies to meet various client requirements in the NetBackup environment. You can perform different operations on the policy like, add, edit, delete, schedule a policy.

Likewise in a protection plan there is a provision to display job operation hierarchy, similarly you need to create an SLP to specify the job operation hierarchy.

Limitations and considerations

Consider the following limitations when creating policies that support cloud workloads:

- Snapshot replication for AWS CSP for IaaS is not supported.

Planning for policies

Policy configuration is flexible enough to meet the various needs of all the Cloud object store accounts in a NetBackup environment. To take advantage of this flexibility, take time to plan before starting to configure the policies.

The following table outlines the steps to take to ensure that you get optimal results from your policy configurations.

Table 1-7 Steps for planning policies

| Step | Action | Description |
|--------|---|--|
| Step 1 | Gather information about the assets that you want to protect. | <p>Gather the following information about each asset:</p> <ul style="list-style-type: none"> ■ The asset names and their region. ■ The approximate number of files for each asset to back up. ■ The typical size of the files. <p>One asset may contain a large amount of data in several files, while the other accounts are smaller with a lesser number of files. To avoid long backup times, include the larger assets in one policy and the smaller ones in another policy. It may be beneficial to create more than one policy for the larger assets.</p> |
| Step 2 | Group the assets based on backup and restore requirements | Divide the different assets into groups according to the different backup and restore requirements. |
| Step 3 | Consider the storage requirements | The storage unit settings apply to all the assets backed up by a policy. If assets have special storage requirements, create separate policies for the assets, even if other factors are the same, such as schedules. |

Table 1-7 Steps for planning policies (*continued*)

| Step | Action | Description |
|--------|---------------------------------|--|
| Step 4 | Consider the backup schedule | <p>Create additional backup policies if the schedules in one policy do not accommodate all the assets that you want to protect.</p> <p>Consider the following factors when deciding to create additional policies:</p> <ul style="list-style-type: none"> ■ Best times for backups to occur. To back up different objects on different schedules may require additional policies with different time schedules. For example, create different policies for night-shift and day-shift backups. ■ How frequently do the assets change? If some assets change more frequently than others, the difference may be enough to warrant creating another policy with a different backup frequency. ■ How long do the backups need to be retained? Each schedule includes a retention setting that determines how long NetBackup keeps the assets that are backed up by the schedule. Because the schedule backs up all the assets that you select for backup, all assets should have similar retention requirements. Do not include the assets whose full backups must be retained forever, together in a policy where full backups are retained for only four weeks. |
| Step 5 | Select exactly what to back up. | You do not need to back up all the discovered assets, unless required. Create queries to select and back up only the required asset(s). |

Creating policies for cloud assets

You can add cloud policies for different types of workloads. Before you create a policy, ensure that you have a Storage lifecycle policy (SLP) created for that policy to view the job operation hierarchy. See [“Adding an SLP”](#) on page 34.

To create a policy:

- 1 On the left, click **Protection > Policies**.
- 2 Click **Add** to create a new policy.
- 3 Configure the settings for the policy.

| | |
|----------------|---|
| Attributes | See “Setting up attributes for IaaS assets” on page 44. See “Setting up attributes for PaaS assets” on page 42. |
| Schedules | See “Creating schedules” on page 46. See “About backup frequency” on page 48. See “About assigning retention periods” on page 49. See “Configuring the Start window” on page 52. See “Configuring the include dates” on page 54. See “Configuring the exclude dates” on page 55. |
| Cloud assets | See “Configuring the cloud assets for IaaS” on page 58. See “Configuring the cloud assets for PaaS” on page 56. |
| Backup options | See “Configuring backup options for IaaS” on page 59. |

Note: This tab appears only for **IaaS** cloud policy type.

- 4 After you are finished configuring the policy, click **Create**.

Setting up attributes for PaaS assets

The **Attributes** tab is used to configure backup settings when you add a new policy or change an existing policy. When you create a policy, you give the policy a name and select a policy type. Not all attributes apply to every policy type, the unavailable attributes are grayed out.

To set the attributes

- 1 On the left, click **Policies**, under **Protection**.
- 2 Enter a name for the policy in the **Policy name** field. For the **Policy type**, select **Cloud**.
- 3 In the **Cloud workload**, select the options **PaaS**.

The **Perform cloud snapshot** option allows snapshot-based protection of cloud-PaaS assets.

- 4 (Optional) If you select the **Perform cloud snapshot** the **Destination** section is not available. This parameter is only applicable for AWS DocumentDB, AWS Neptune, RDS Custom Oracle, RDS Custom SQL, and Redshift cluster assets using full backup schedule.
- 5 In the **Destination** section, configure the following data storage parameters:
 - The **Data classification** attribute specifies the classification of the storage lifecycle policy that stores the backup. For example, a backup with a gold classification must go to a storage unit with a gold data classification. By default, NetBackup provides four data classifications: Platinum, Gold, Silver, and Bronze. This attribute is optional and applies only when the backup is to be written to a storage lifecycle policy. If you select **No data classification**, the policy uses the storage selection that is displayed in the **Policy storage** list. If a data classification is selected, all the images that the policy creates are tagged with the classification ID.
 - The **Policy storage** attribute specifies the storage destination for the policy's data. You can override these selections from the **Schedule** tab.
- 6 The **Limit jobs per policy** attribute limits the number of jobs that NetBackup performs concurrently when the policy is run. By default, the box is cleared, and NetBackup performs an unlimited number of backup jobs concurrently. Other resource settings can limit the number of jobs.

A configuration can contain enough devices so that the number of concurrent backups affects performance. To specify a lower limit, select the box and specify a value from 1 to 999.
- 7 In the **Job priority** field, enter a value from 0 to 99999. This number specifies the priority that a policy has as it competes with other policies for resources. The higher the number, the greater the priority of the job. NetBackup assigns the first available resource to the policy with the highest priority.

- 8 To activate the policy, select the option **Go into effect at**, and set the date and time of activation. For example, if today is a Monday and you enter Wednesday at 12:00 A.M., the policy does not run until that time or later. The policy must be active for NetBackup to use it.

To deactivate a policy, clear the option. Inactive policies are available in the **Policies** list. To resume backups, select this option again. Make sure that the date and time are set to the time that you want to resume backups.

Use this attribute to configure a series of policies in advance of when the policies need to become active.

- 9 The **Keyword phrase** attribute is a phrase that NetBackup associates with all backups or archives based on the policy. Only the Windows and UNIX client interfaces support keyword phrases.

You can use the same keyword phrase for more than one policy. The same phrase for multiple policies makes it possible to link backups from related policies. For example, use the keyword phrase "legal department documents" for backups of multiple clients that require separate policies, but contain similar types of data.

The phrase can be a maximum of 128 characters in length. All printable characters are permitted, including spaces and periods. By default, the keyword phrase is blank.

Setting up attributes for IaaS assets

The **Attributes** tab is used to configure backup settings when you add a new policy or change an existing policy. When you create a policy, you give the policy a name and select a policy type. Not all attributes apply to every policy type, the unavailable attributes are grayed out.

For IaaS cloud, you need to create a Storage Lifecycle Policy (SLP) with Snapshot as primary operation and backup from snapshot as secondary operation along with the storage unit. See ["About storage lifecycle policies"](#) on page 33.

To set attributes

- 1 On the left, click **Policies**, under **Protection**.
- 2 Enter a name for the policy in the **Policy name** field. For the **Policy type**, select **Cloud**.
- 3 Under **Cloud workload**, select the options **IaaS**.
- 4 In the **Destination** section, configure the following data storage parameters:
 - The **Data classification** attribute specifies the classification of the storage lifecycle policy that stores the backup. For example, a backup with a gold

classification must go to a storage unit with a gold data classification. By default, NetBackup provides four data classifications: Platinum, Gold, Silver, and Bronze. This attribute is optional and applies only when the backup is to be written to a storage lifecycle policy. If you select **No data classification**, the policy uses the storage selection that is displayed in the **Policy storage** list. If a data classification is selected, all the images that the policy creates are tagged with the classification ID.

- The **Policy storage** attribute specifies the storage destination for the policy's data. You can also create an SLP and select it from the the dropdown. You can override these selections from the **Schedule** tab.

- 5 The **Limit jobs per policy** attribute limits the number of jobs that NetBackup performs concurrently when the policy is run. By default, the box is cleared, and NetBackup performs an unlimited number of backup jobs concurrently. Other resource settings can limit the number of jobs.

A configuration can contain enough devices so that the number of concurrent backups affects performance. To specify a lower limit, select the box and specify a value from 1 to 999.

- 6 In the **Job priority** field, enter a value from 0 to 99999. This number specifies the priority that a policy has as it competes with other policies for resources. The higher the number, the greater the priority of the job. NetBackup assigns the first available resource to the policy with the highest priority.

- 7 To activate the policy, select the option **Go into effect at**, and set the date and time of activation. For example, if today is a Monday and you enter Wednesday at 12:00 A.M., the policy does not run until that time or later. The policy must be active for NetBackup to use it.

To deactivate a policy, clear the option. Inactive policies are available in the **Policies** list. To resume backups, select this option again. Make sure that the date and time are set to the time that you want to resume backups.

Use this attribute to configure a series of policies in advance of when the policies need to become active.

- 8 The **Keyword phrase** attribute is a phrase that NetBackup associates with all backups or archives based on the policy. Only the Windows and UNIX client interfaces support keyword phrases.

You can use the same keyword phrase for more than one policy. The same phrase for multiple policies makes it possible to link backups from related policies. For example, use the keyword phrase "legal department documents" for backups of multiple clients that require separate policies, but contain similar types of data. The phrase can be a maximum of 128 characters in length. All printable characters are permitted, including spaces and periods. By default, the keyword phrase is blank.

Creating schedules

The schedules defined on the **Schedules** tab determine when the backups occur for the selected policy. Each schedule also includes various criteria, such as how long to retain the backups.

Schedule attributes appear on the following tabs:

| | |
|--------------------------|--|
| Attributes tab | Schedule the time and frequency at which a task runs, along with other scheduled attributes. |
| Start Window tab | Schedule the time of each day that a task runs. |
| Exclude Days tab | Indicate the days that a job cannot run. |
| Include Dates tab | Schedule the run days for a task by indicating specific dates, recurring weekdays, recurring days of the month. (This tab appears only when Calendar is selected as the Schedule type .) |

To create a schedule for a policy

- 1 On the left, click **Policies**, under **Protection**. Click the **Schedules** tab. Under **Backup schedules**, click **Add**. Click the **Attributes** tab.
- 2 In the Attributes tab, enter a name for the schedule in the **Name** field.
- 3 Select the **Type of backup**. For IaaS workloads, only **Full backup** is supported.
 - **Full backup** - Backs up all of the files that are specified in the policy. The files are backed up, regardless of when the files were last modified or backed up. Full backups occur automatically according to schedule criteria. If you run incremental backups, you must also schedule a full backup to perform a complete restore.
 - **Differential incremental backup** - Backs up the files that changed since the last successful incremental (differential or cumulative) or full backup. All files are backed up if no previous backup was done. Differential incremental backups occur automatically according to schedule criteria. A complete restore requires the last full backup and all differential incremental backups that occurred since the last full backup.
 - **Archived redo log backup** - In this method, NetBackup backs up the data that is changed since the subsequent full or incremental backup. Archive backups reduce the full and incremental backup windows significantly. See [“About archive redo log backup for PaaS workloads”](#) on page 140.

Note: Amazon (AWS) RDS Oracle supports transaction log archiving.

- 4 Under the **Destination**, the appropriate parameters are visible:
 - **Override policy storage selection** attribute works as follows:
 - **Enabled:** Instructs the schedule to override the Policy storage as specified on the policy **Attributes** tab. Select the storage from the list of previously configured storage units and storage lifecycle policies. If the list is empty, no storage is configured.
 - **Disabled:** Instructs the schedule to use the Policy storage as specified on the policy **Attributes** tab.
- 5 Under **Schedule** type, select **Calendar** or **Frequency**.
 - **Calendar** Calendar-based schedules let you create a job schedule based on a calendar view. Select **Calendar** to display the **Include dates** tab. Enable Retries allowed after run day to have NetBackup attempt to complete the schedule until the backup is successful. With this attribute enabled, the schedule attempts to run, even after a specified run day has passed.

- **Frequency** Use the **Frequency** attribute to specify how much time must elapse between the successful completion of a scheduled task and the next attempt. For example, assume that a schedule is set up for a full backup with a frequency of one week. If NetBackup successfully completes a full backup for all clients on Monday, it does not attempt another backup for this schedule until the following Monday. To set the frequency, select a frequency value from the list. The frequency can be seconds, minutes, hours, days, or weeks.
- 6 Specify a **Retention** period for the backups. This attribute specifies how long NetBackup retains the backups. To set the retention period, select a period (or level) from the list. When the retention period expires, NetBackup deletes information about the expired backup. After the backup expires, the objects in the backup are unavailable for restores. For example, if the retention is 2 weeks, data can be restored from a backup that this schedule performs for only 2 weeks after the backup.
 - 7 Click **Add** to add the attributes, or click **Add and add another** to add a different set of attributes for another schedule.

About backup frequency

To determine backup frequency, consider how often data changes. For example, determine if files change several times a day, once a day, weekly, or monthly.

Typically, sites perform daily backups to preserve daily work. Daily backups ensure that only one day's work is lost in case of a disk failure. More frequent backups are necessary when important data changes many times during the day and the changes would be difficult to reconstruct.

Daily backups are usually the differential incremental backups that record the changes since the last differential incremental or full backup. Differential incremental backups conserve resources because they use less storage and take less time to perform than full backups.

Full backups usually occur less frequently than differential incremental backups but should occur often enough to avoid accumulating consecutive differential incremental backups. A large number of differential incremental backups between full backups increases the time it takes to restore a file. The time increases because of the effort that is required to merge the differential incremental backups when files and directories upon restore.

Consider the following when setting the frequency for full backups:

- Extend the time between full backups for the files that seldom change. A longer frequency uses fewer system resources. It also does not significantly increase

recovery time because the differential incremental backups between full backups are smaller.

- Decrease the time between full backups for the files that change frequently. A shorter frequency decreases restore time. A shorter time between full backups can also use fewer resources. It reduces the cumulative effect of the longer differential incremental backups that are necessary to keep up with frequent changes in the files.

To achieve the most efficient use of resources, ensure that most of the files in a given policy change at about the same rate. For example, assume that half of the files in a policy selection list change frequently enough to require a full backup every week. However, the remaining files seldom change and require monthly full backups only. If all the files are in the same policy, full backups are performed weekly on all the files. This wastes system resources because half the files need full backups only once a month. A better approach is to divide the backups into two policies, each with the appropriate backup schedule, or to use synthetic backups.

If more than one automatic schedule is due for a client within a policy, the backup frequency determines the schedule that NetBackup uses as follows:

- Jobs from the schedule with the lower frequency (longer period between backups) always have higher priority. For example, a schedule that has a backup frequency of one month takes priority over a schedule with a backup frequency of 2 weeks.
- When two schedules are each due to run, the schedule with the schedule name that is first in alphabetical order runs first. Alphabetical priority occurs if both of the following are true:
 - Each schedule is within the defined time window.
 - Each schedule is configured with the same frequency value.

NetBackup prioritizes the example schedules in the following order:

Table 1-8 Examples of schedule frequency and priority

| Schedule Name | Frequency | Priority |
|--------------------------------|-----------|----------|
| monthly_full | One month | First |
| weekly_full | One week | Second |
| daily_differential_incremental | One day | Third |

About assigning retention periods

The retention period for data depends on the likelihood of restoring information from media after a certain period of time. Some types of data (financial records, for

example) have legal requirements that determine the retention level. Other data (preliminary documents, for example) can probably be expired when the final version is complete.

A backup’s retention also depends on what needs to be recovered from the backup. For example, if day-to-day changes are critical, keep all the incremental backups in addition to the full backups for as long as the data is needed. If incremental backups only track work in progress toward monthly reports, expire the incremental backups sooner. Rely on the full backups for long-term recovery.

Establish some guidelines that apply to most of the data to determine retention periods. Note the files or the directories that have retention requirements outside of these guidelines. Plan to create separate policies for the data that falls outside of the retention requirement guidelines. For example, place the files and directories with longer retention requirements in a separate policy. Schedule longer retention times for the separate policies without keeping all policies for the longer retention period.

The following table describes recommended retention periods for different types of backups.

Table 1-9 Recommended retention periods for different types of backups

| Type of backup | Description |
|--|---|
| Full backup | Specify a time period that is longer than the frequency setting for the schedule. (The frequency is how often the backup runs). For example, if the frequency is one week, specify a retention period of 2-4 weeks. Two to 4 weeks provides enough of a margin to ensure that the current full backup does not expire before the next full backup occurs. |
| Differential incremental backup | Specify a time period that is longer than the period between full backups. For example, if full backups occur weekly, save the incremental backups for 2 weeks. |

The following table suggests several ways that you can prevent backups from expiring earlier than desired.

Table 1-10 Suggestions for preventing prematurely expired backups

| Item | Description |
|--------------------------------------|---|
| Retention period | <p>Assign an adequate retention period. NetBackup does not track backups after the retention period expires. Recovering files is difficult or impossible after the retention period expires.</p> <p>For the backups that must be kept for more than one year, set the retention period to infinite.</p> |
| Full backups and incremental backups | <p>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.</p> |
| Archived schedules | <p>Set the retention period to infinite.</p> |
| Tape | <p>Set the retention period to infinite. If infinite is unacceptable because of NetBackup database space limitations, set the retention period to match the length of time that the data is to be retained.</p> |

Another consideration for data retention is off-site storage of the backup media. Off-site storage protects against the disasters that may occur at the primary site. Consider the following off-site storage methods as precautions for disaster recovery:

- Use the duplication feature to make a second copy for off-site storage.
- Send monthly or weekly full backups to an off-site storage facility.
 To restore the data, request the media from the facility. To restore a total directory or disk with incremental backups requires the last full backup plus all incremental backups.
- Configure an extra set of schedules to create the backups to use as duplicates for off-site storage.

Regardless of the method that is used for off-site storage, ensure that adequate retention periods are configured.

By default, NetBackup stores each backup on a tape volume that contains existing backups at the same retention level. If a backup has a retention level of 2, NetBackup stores it on a tape volume with other backups at retention level 2. When NetBackup encounters a backup with a different retention level, it switches to an appropriate volume. Because tape volumes remain assigned to NetBackup until all the backups on the tape expire, this approach results in more efficient use of media. One small backup with an infinite retention prevents a volume from being reused, even if all other backups on the volume expired.

If you keep only one retention level on each volume, do not use any more retention levels than necessary. Multiple retention levels increase the number of required volumes.

Note: Retention levels can be mixed on disk volumes with no restrictions.

Configuring the Start window

The **Start window** tab provides controls for setting periods during which NetBackup can start jobs using a schedule. Periods are referred to as windows. Configure the windows to satisfy the requirements necessary to complete a job.

You can also perform other operations on the schedule such as - Delete, Clear, Duplicate, and Undo.

To configure the Start window:

1 Click the **Schedules** tab. Under **Backup schedules**, click **Add**. Click the **Start window** tab.

2 To indicate the opening of the time window, do the following:

Drag your cursor in the time table.

Click the day and time when you like the window to start and drag it to the day and time when you like the window to close.

Use the settings in the dialog box.

- In the **Start day** field, select the first day that the window opens.
- In the **Start time** field, select the time that the window opens.

3 To indicate the closing of the time window, do one of the following:

Drag your cursor in the time table.

Click the day and time when you like the window to start and drag it to the day and time when you like the window to close.

Enter the duration of the time window.

Enter a length of time in the **Duration (days, hours, minutes)** field.

Indicate the end of the time window.

- Select a day in the **End day** list.
- Select a time in the **End time** field.

Time windows show as bars in the schedule display.

Specify enough time to allow all clients in the policy to complete a backup.

Consider allowing extra time in the schedule in case the schedule starts late due to factors outside of NetBackup. (Delays due to unavailable devices, for example.) Otherwise, all backups may not have a chance to start.

4 As necessary, do any of the following:

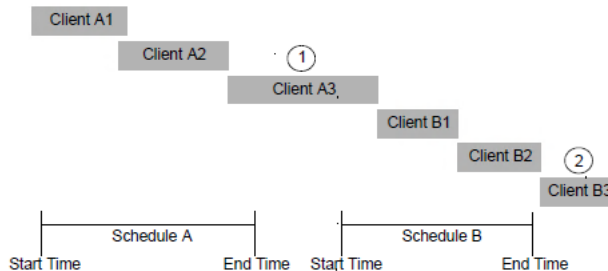
- Click **Delete**. Deletes the selected time window.
- Click **Clear**. Deletes all the time windows from the schedule display.
- Click **Duplicate**. Replicates the time window for the entire week.
- Click **Undo**. Erases the last action.

5 Do one of the following:

- Click **Add**. To save the time window and leave the dialog box open.
- Click **Add and Add another**. To save the time window and add another.

Example of schedule duration

This example illustrates the effect of schedule duration on two full backup schedules. The start time for Schedule B begins shortly after the end time for the previous Schedule A. Both schedules have three clients with backups due.



The image illustrates the following points:

- Point 1 Client A3 starts within the Schedule A time window but does not complete until after the Schedule B start time. However, Client A3 runs to completion even if the window closes while the backup is running. Client B1, on Schedule B, begins as soon as Client A3 completes.
- Point2 Schedule A does not leave enough time for all the clients on Schedule B to be backed up. Consequently, Client B3 is unable to start because the time window has closed. Client B3 must wait until the next time NetBackup runs Schedule B.

Configuring the include dates

The **Include dates** tab appears in the Add schedule or Edit schedule tabs. For the tab to display, you must select the **Calendar** option as the **Schedule type** on the **Attributes** tab.

The tab displays a calendar of three consecutive months. Use the lists at the top of the calendar to change the first month or year displayed.

Use the **Calendar** option in the policy **Attributes** tab to create a job schedule based on a calendar view. The **Include dates** tab lets you configure the schedules that run according to specific days, on recurring week days, or on recurring days of the month.

Note: Using the calendar schedule, if a green circle does not appear on a day, the day is not included in the schedule.

If **Retries allowed after run day** is enabled, a job may run on a day that is not included in the schedule.

When a new calendar schedule is created with **Retries allowed after run day** enabled, the schedule runs its first job on the next day when the backup window is open. That day may be before the first run day that is included in the schedule.

To use a calendar to schedule run days:

- 1 In the **Attributes** tab, enable the **Calendar** option.
- 2 Select the **Include dates** tab.
- 3 Use one or more methods to schedule the days on which jobs can run:
 - Select the day(s) on the three-month calendar on which you want the jobs to run. Use the drop-down lists at the top of the calendar to change the months or year.
 - To indicate **Recurring week days**:
 - Click **Set all** to select all of the days in every month for every year.

- Click **Clear all** to remove all existing selections.
 - Select a box in the matrix to select a specific day to include for every month.
 - Click the column head of a day of the week to include that day every month.
 - Click the **1st, 2nd, 3rd, 4th,** or **Last** row label to include that week every month.
 - To indicate **Recurring days of the month:**
 - Click **Set all** to select all of the days in every month.
 - Click **Clear all** to remove all existing selections.
 - Check a box in the matrix to select that day to include each month.
 - Click **Last** to include the last day of every month.
 - To indicate **Specific dates:**
 - Click **New**. Enter the month, day, and year in the dialog. The date appears in the **Specific dates** list.
 - To delete a date, select the date in the list. Click **Delete**.
- 4 Click **Add** to save the included days.

Configuring the exclude dates

Use the **Exclude dates** tab to exclude specific days from a schedule for a backup policy. If a day is excluded from a schedule, jobs do not run on that day. The tab displays a calendar of three consecutive months. Use the lists at the top of the calendar to change the first month or year displayed.

To exclude a day from a schedule:

- 1 On the left, click **Policies**, under **Protection**. Click the **Schedules** tab. Under **Backup schedules**, click **Add**. Click the **Exclude dates** tab.
- 2 Use one or more methods to indicate the days to exclude:
 - Select one or more days on the 3-month calendar that you want to exclude. Use the dropdown list at the top of the calendar to change the months or years.
 - To indicate **Recurring week days:**
 - Click **Set all** to select all of the days in every month for every year.
 - Click **Clear all** to remove all existing selections.

- Select a box in the matrix to select a specific day to exclude for every month.
 - Click the column head of a day of the week to exclude that day every month.
 - Click the **1st, 2nd, 3rd, 4th, or Last** row label to exclude that week every month.
 - To indicate **Recurring days of the month:**
 - Click **Set all** to select all of the days in every month.
 - Click **Clear all** to remove all existing selections.
 - Select a box in the matrix to select that day to exclude each month.
 - Click **Last** to exclude the last day of every month.
 - To indicate **Specific dates:**
 - Click **New**. Enter the month, day, and year in the dialog box. The date appears in the **Specific dates** list.
 - To delete a date, select the date in the list. Click **Delete**.
- 3** Click **Add** to save the changes.

Configuring the cloud assets for PaaS

The **Cloud assets** tab lets you configure auto-managed database in the cloud environment.

The multiple copies option to add a schedule to the PaaS assets is not supported.

Table 1-11 Cloud asset type PaaS against the cloud provider

| Cloud Provider | DB service |
|-----------------------|---|
| Amazon Web Services | Aurora MySQL Aurora PostgreSQL DynamoDB MariaDB MySQL Oracle PostgreSQL Redshift SQL Server Custom SQL Custom Oracle DocumentDB Neptune |
| Microsoft Azure | Cosmos DB for MongoDB Cosmos DB for noSQL MariaDB MySQL PostgreSQL SQL Managed Instance SQL Server |
| Google Cloud Platform | GCP SQL Server GCP MySQL GCP PostgreSQL GCP BigQuery |

To add asset to the policy

- 1** On the **Cloud assets** tab, select the **Provider** from the dropdown.
- 2** Select the asset type from the **DB service** dropdown.
- 3** Click on the **Add assets**.

- 4 The **Add assets** pane displays the asset type selected in step 2. Select a single or a multiple asset type.

You can also add single or multiple intelligent groups from the **Add asset** pane.

- 5 Click **Add**. The asset type is added in the list on the **Cloud asset** tab.

To remove the asset

- 1 On the **Cloud assets** tab, you can remove the asset type from the list.
- 2 Select the check box against the asset type and click **Remove**. You can also use the remove option from the **Actions > Remove**.

Configuring the cloud assets for IaaS

The **Cloud assets** tab lets you configure assets like virtual machines, applications, and volumes in the cloud. You can also select existing intelligent groups to configure in the cloud environment.

The cloud assets for backup differ as per the cloud provider. They are:

Table 1-12 Cloud asset type IaaS against the cloud provider

| Cloud provider | Cloud assets for backup |
|-----------------------------|--|
| Amazon Web Services | Virtual machine Application Volume |
| Google Cloud Provider | Virtual machine Application Volume |
| Microsoft Azure | Virtual machine Application Volume |
| Microsoft Azure Stack Hub | Virtual machine Application Volume |
| Oracle Cloud Infrastructure | Virtual machine Oracle applications |

To add asset to the policy

- 1 On the **Cloud assets** tab, select the **Provider** from the dropdown.
- 2 Select the asset type from the **Assets for backup** dropdown.
- 3 Click on the **Add assets**.
- 4 The **Add assets** pane displays the asset type (virtual machine, application, or volume) selected in step 2. Select a single or a multiple asset type.
 Select one or more assets or intelligent groups from the **Add asset** pane.
- 5 Click **Add**. The asset type is added in the list on the **Cloud asset** tab.

You can remove the asset using the **Remove** option.

To remove the asset

- 1 On the **Cloud assets** tab, you can remove the asset type (virtual machine, application, or volume) from the list.
- 2 Select the check box against the asset type and click **Remove**. You can also use the remove option from the **Actions > Remove**.

Configuring backup options for IaaS

The **Backup options** tab includes multiple options for enabling backups for recovery of files and folders and other options. This tab is applicable only for IaaS cloud type.

The options in the **Backup options** tab differ as per the cloud service provider that is selected in the **Cloud assets** tab.

Table 1-13 Backup options against the cloud provider

| Cloud Provider | Backup options |
|-----------------------|--|
| Amazon Web Services | Enable granular recovery for files or folders. Exclude selected disks from backups. |
| Google Cloud Platform | Enable granular recovery for files or folders. Enable a regional snapshot. Exclude selected disks from backups. |
| Microsoft Azure | Enable granular recovery for files or folders. Specify a snapshot destination resource group. Exclude selected disks from backups. |

Table 1-13 Backup options against the cloud provider (*continued*)

| Cloud Provider | Backup options |
|-----------------------------|--|
| Microsoft Azure Stack Hub | Enable granular recovery for files or folders. Specify a snapshot destination resource group. Exclude selected disks from backups. |
| Oracle Cloud Infrastructure | Enable granular recovery for files or folders. |

Managing cloud policies

You can perform multiple operations on the cloud policy using the NetBackup web UI.

Table 1-14 Operations on cloud policy

| Operations | Description |
|-----------------------|---|
| Edit | Except for the name of the policy, all the attributes are editable. |
| Copy policy | You can create a copy of the policy. The new copy opens in edit mode. |
| Delete | You can delete the policy using this option. |
| Activate / Deactivate | You can activate or deactivate the policy. |
| Manual backup | You can initiate the manual backup of the policy. Manual backups are possible only for active policy. |

- See [“Copy a policy”](#) on page 60.
- See [“Deactivating or deleting a policy”](#) on page 61.
- See [“Manually backup assets”](#) on page 62.

Copy a policy

Copying a policy lets you reuse similar policy attributes, schedules, and cloud objects among your policies. You can also reuse complex queries by copying policies, to save time.

To copy a policy:

- 1 On the left, click **Policies**. All the policies that you have the privilege to view are displayed in the **Policies** tab.
- 2 Click the ellipsis menu (three dots) in the row of the policy that you want to copy. Click **Copy policy**.

Alternatively, select the option in the row of the policy, click **Copy policy** at the top of the table.
- 3 In the **Copy policy** dialog, optionally, change the name of the policy in the **Policy to copy** field.
- 4 Enter the name of the new policy, in the **New policy** field.
- 5 Click **Copy** to initiate copying.

Deactivating or deleting a policy

Deactivating a policy has the following implications:

- You cannot perform manual backups for deactivated policies.
- Scheduled backups in the deactivated policies are not triggered.
- Operations such as edit, copy, and delete works normally.
- Copying the deactivated policy creates a new policy in the deactivated state.

When you delete a policy, the scheduled backups, which were configured in that policy, are not conducted.

To deactivate or delete a policy:

- 1 On the left, click **Policies**. All the policies that you have the privilege to view are displayed in the **Policies** tab.
- 2 Click the ellipsis menu (three dots) in the row of the policy that you want to copy. Click Deactivate or **Delete** as required.

Alternatively, select the option in the row of the policy, click **Deactivate** or **Delete** as required, at the top of the table.

The policies get deactivated immediately. To reactivate the policy again, click the ellipsis menu (three dots) in the row of the deactivated policy and click **Activate**.

- 3 If you delete a policy, click **Delete** in the confirmation box.

Manually backup assets

Apart from the scheduled backups performed by the policies, you can perform ad hoc, manual backups for a policy as required.

To perform a manual backup:

- 1 On the left, click **Policies**. All the policies that you have the privilege to view are displayed in the **Policies** tab.
- 2 Click the ellipsis menu (three dots) in the row of the policy for which you want to perform backup. Click **Manual backup**.

Alternatively, select the option in the row of the policy, and click **Manual backup**, at the top of the table.
- 3 In the **Manual backup** dialog, select the schedule that you want to use for the backup. You can see the schedules defined in the policy.
- 4 Select one or more clients you want to back up. If you do not select any, all clients are backed up.
- 5 Click **OK** to start the backup.

Scan for malware

NetBackup provides support for scanning Cloud assets for malware through the Cloud workload type.

For triggering malware scan, the scan host must be configured. For more information on configuring the scan host, refer to the 'Scan host configurations' chapter in *NetBackup Security and Encryption Guide*.

Scanning backup images

This section describes the procedure for scanning client backup images of a particular policy for malware.

To scan policy of client backup images for malware

- 1 On the left, select **Detection and reporting > Malware detection**.
- 2 On the **Malware detection** page, select **Scan for malware**.
- 3 In the **Search by** option, select **Backup images**.

Select one of the following scan types:

- Malware scan - Select this option to scan images using default malware scan.
- YARA scan - Select this option to scan images using YARA rules.

Click the **Select threat feeds** option.

On the **Select threat feeds for scanning** dialog box, select the required YARA rules or a zip file of YARA rules that you have uploaded earlier.

All the following steps are applicable for the scan type: Malware scan.

- 4 In the **Scanner host pool** option, search and select the appropriate host pool name from the list of scanner host pools listed in **Select malware scanner host pool**.

Note: Scan host from the selected scan host pool must be able to access the instant access mount created on the storage server which is configured with NFS/SMB share type.

- 5 In the search criteria, review and edit the following:
 - **Policy name**
Only supported policy types are listed.
 - **Client name**
Displays the clients that have backup images for a supported policy type.
 - **Policy type**
Displays all the supported policies which are enabled for malware scanning.
 - **Type of backup**
 - **Copies**
If the selected copy does not support instant access, then the backup image is skipped for the malware scan.
 - **Disk pool**
MSDP (PureDisk), OST (for example, Data Domain) and AdvancedDisk storage type disk pools are listed.
 - **Disk type**
MSDP (PureDisk), OST (for example, Data Domain) and AdvancedDisk disk types are listed.
 - **Infection status**
The malware infected status of the backup images can be searched based on the following types: infection detected by malware scan, file hash search, not infected, not scanned or all.
 - For the **Select the timeframe of backups**, verify the date and the time range or update it.

- On selecting the **Abort malware scan on detecting an infection** option, clean recovery would not be supported for infected images.
- 6 Click **Search**.
 - 7 Select the search criteria and ensure that the selected scan host is active and available.
 - 8 From the **Select the backups to scan** table select one or more images for scan.
 - 9 Click **Scan for malware**.
 - 10 After the scan is initiated, the **Scan status** is displayed.

The following are the status fields:

- **Not scanned**
- **Not infected**
- **Infected**
- **Failed**
- **In progress**
- **Pending**

Hover over the status to view the reason for the failed scan.

Note: You can cancel the malware scan for one or more in progress and pending jobs.

- **Infected - Malware scan aborted**

Assets by workload type

This section describes the procedure for scanning Cloud VM assets for malware.

Note: For YARA scanning, only Kubernetes is supported.

To scan the supported assets for malware, perform the following:

- 1 On the left, select the supported workload under **Workloads**.
- 2 Select the resource which has backups completed.
 For example, Cloud VM.
- 3 Select **Actions > Scan for malware**.

- 4 On the **Malware scan** page, perform the following:
 - Select the date range for the scan by selecting **Start date/time** and **End date/time**.
 - Select **Scanner host pool**.
 - From the **Current infection status** list select one of the following:
 - **Not scanned**
 - **Not infected**
 - **Infection detected by malware scan**
 - **Infection detected by file hash search**
 - **All**

- 5 Click **Scan for malware**.

Note: The malware scanner host can initiate a scan of three images at the same time.

- 6 After the scan starts, you can see the **Scan status** on **Malware detection**, the following fields are visible:
 - **Not scanned**
 - **Not infected**
 - **Infected**
 - **Failed**

Note: Any backup images that fail validation are ignored.

- **In progress**
- **Pending**

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 1-15 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 | NetBackup creates new snapshots and associates the snapshots to the peer resource group if the following conditions are met: <ul style="list-style-type: none"> ■ The peer resource group is created. ■ The peer resource group is in the same region as the resource group. If the conditions are not met, snapshot jobs fail. |

Table 1-15 Configurations and outcome (*continued*)

| Resource group prefix | Protect assets even if prefixed Resource Groups are not found check box | Outcome |
|-----------------------|---|---|
| Specified | Selected | <p>NetBackup creates new snapshots and associates the snapshots to the peer resource group if the following conditions are met:</p> <ul style="list-style-type: none"> ■ The peer resource group is created. ■ The peer resource group is in the same region as the resource group. <p>If a peer resource group is not created or is in a different region then the newly created snapshot is associated 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 1-16 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. | <p>Snapshots are created in the peer resource group.</p> |

Table 1-16 Example configurations (*continued*)

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

Table 1-16 Example configurations (*continued*)

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

Troubleshoot resource group permissions

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

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.

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

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

Note: Accelerator is not supported for Google Cloud Platform (GCP), and Oracle Private Cloud Appliance.

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 Hub 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.

Considerations for using Accelerator with OCI

Consider the following when using Accelerator with OCI backups.

- If a policy or protection plan do not have accelerator enabled at the time of creation, you cannot enable accelerator for policy and protection plan by editing that policy. Similarly, if you upgrade NetBackup from an older version that did not support Accelerator, you cannot enable Accelerator for the policies or protection plans created in the older version. In both the cases, you must create a new policy to use Accelerator.
- If you disable Accelerator for a policy using the `bp.conf` file, you cannot enable Accelerator for that policy again. You must create a new policy.
- Accelerator backups are disabled if the volumes are resized between backups. NetBackup runs a full backup.
- OCI volumes of `Lower Cost performance level (0 VPU/GB)` type are not supported for Accelerator backups.

- For Accelerator backups, there must be no full snapshots between two consecutive snapshots. The first snapshot can be a full snapshot or an incremental snapshot. The second (current) snapshot must be an incremental snapshot. If both snapshots are incremental, no full snapshot must occur between the two incremental backups.

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 using protection plan

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 Hub, 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.

Configuring backup schedules for cloud workloads using protection plan

- 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 [“Protecting PaaS assets in AWS” on page 138](#) on page 138.

- 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.

Configuring backup schedules for cloud workloads using protection plan

- 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 80.
 - 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, OCI, 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 GCP

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

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

Snapshot destination resource group for Azure and Azure Stack Hub

You can choose to specify a snapshot destination peer resource group while creating a protection plan for Azure or Azure Stack Hub. While the previous functionality of

defining a peer resource group by specifying a prefix still exists, you can now directly associate a snapshot to an existing peer resource group at the time of creating a protection plan.

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

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

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

Excluding selected disks from 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.

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.

AWS 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**

In case of replication to a different region in a different AWS account, the source region (where the VM to be protected exists) must be enabled in the target AWS accounts.

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 15.
- 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 using protection plan”](#) on page 74.
- 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 24.
- 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 152.

- 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 24.
- 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 1-17 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 1-18 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 1-19 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.
- For alternate location restore of Oracle databases, to an alternate VMs that is a clone of the original VM, remove the fstab entries and restart the daemon, using the command:

```
# systemctl daemon-reload
```

Before you begin

Ensure that the database is prepared for snapshots. For details review the plug-in configuration notes in the *NetBackup Snapshot Manager for Cloud Install and Upgrade Guide*.

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 *NetBackup Web UI 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.

For more information on consideration and prerequisites required for preparing the VMs, refer to the following section:

See [“Recovering AWS or Azure VMs to VMware”](#) on page 173.

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/opensv/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.

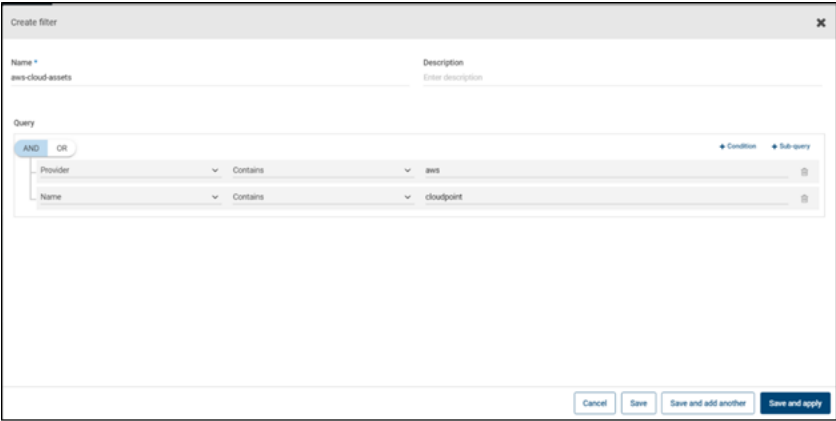
- 3 Click the **Create filter** option to define a custom filter based on attributes to list assets. in the Virtual machines, Applications, PaaS, or Volumes tab.
- 4 To create a filter, enter the details for the following parameters:

| 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 28.

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



This example uses **AND** to narrow the scope of the query: It selects only the 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 1-20 Query keywords

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

Table 1-20 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 1-21 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. |

Protecting PaaS assets

This chapter includes the following topics:

- [Protecting PaaS assets](#)
- [Steps to protect PaaS assets](#)
- [Prerequisites for protecting PaaS assets](#)
- [Enabling binary logging for MySQL and MariaDB databases](#)
- [Enabling backup and restore in Kubernetes](#)
- [Prerequisites for protecting Amazon RDS SQL Server database assets](#)
- [Protecting RDS Custom instances](#)
- [Protecting Azure Managed Instance databases](#)
- [Prerequisites for database level discovery](#)
- [Limitation and considerations](#)
- [Installing the native client utilities](#)
- [Configuring storage for different deployments](#)
- [Configuring the storage server for instant access](#)
- [About incremental backup for PaaS workloads](#)
- [Configuring incremental backups for Azure MySQL server](#)
- [About archive redo log backup for PaaS workloads](#)
- [About Auto Image Replication for PaaS workloads](#)
- [Discovering PaaS assets](#)

- [Viewing PaaS assets](#)
- [Managing PaaS credentials](#)
- [Add protection to PaaS assets](#)

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.

Steps to protect PaaS assets

This section gives you the sequence of steps that you need to perform to get started protecting your PaaS workload on different cloud platforms.

Cloud platforms:

- See [“Protecting PaaS assets in AWS”](#) on page 94.
- See [“Protecting PaaS assets in Azure”](#) on page 97.
- See [“Protecting PaaS assets in GCP”](#) on page 100.

Protecting PaaS assets in AWS

Follow these steps to protect your PaaS assets in AWS. The "Applicable to?" column shows for which database or environment a particular step is required. The Steps column gives you the links to the steps that you need to perform. Ensure that you perform each step completely before you proceed to the next step.

The steps that are applicable to all databases are mandatory steps. You need to perform the database-specific steps only when you protect that database.

Table 2-1 Steps to protect PaaS assets in AWS

| Step no | Applicable to? | Steps |
|---------|-------------------------------|---|
| 1 | All databases | Review the prerequisites for NetBackup PaaS protection. See “Prerequisites for protecting PaaS assets” on page 103. |
| | RDS SQL Server databases | For RDS SQL Server databases, consider these prerequisites. See “Prerequisites for protecting Amazon RDS SQL Server database assets” on page 107. |
| | RDS Custom instance databases | For RDS Custom instance databases, consider the following. See “Protecting RDS Custom instances” on page 108. |
| 2 | All databases | Consider the limitations and considerations of NetBackup protection for the databases that you want to protect. See “Limitation and considerations” on page 112. |
| 3 | All databases | Install the native client utilities. Each native client utility supports a set of databases. You can find the list of supported databases for each utility in the relevant sections. See “Installing the native client utilities” on page 129. |
| 4 | All databases | Configure storage for different NetBackup deployments. See “Configuring storage for different deployments” on page 136. |
| 5 | All databases | Configure the storage server for instant access. See “Configuring the storage server for instant access” on page 137. |

Table 2-1 Steps to protect PaaS assets in AWS (*continued*)

| Step no | Applicable to? | Steps |
|---------|--|--|
| 6 | All databases | <p>NetBackup uses the NetBackup Snapshot Manager component to provide protection to the PaaS databases. Make sure that you do the following:</p> <ul style="list-style-type: none"> ■ Install and configure NetBackup Snapshot Manager. ■ Install and configure the AWS plug-in for NetBackup Snapshot Manager. ■ Configure the required permissions for AWS in the NetBackup Snapshot Manager. <p>See the <i>NetBackup™ Snapshot Manager for Cloud Install and Upgrade Guide</i> for more details.</p> |
| 7 | When NetBackup is deployed in an EKS cluster | <p>Enable backup and restore in EKS.</p> <p>See “Enabling backup and restore in Kubernetes” on page 106.</p> |
| 8 | MySQL and MariaDB databases | <p>Enable binary logging.</p> <p>See “Enabling binary logging for MySQL and MariaDB databases” on page 106.</p> |
| 9 | All databases | <p>Consider Auto Image Replication (AIR) for extra protection.</p> <p>See “About Auto Image Replication for PaaS workloads” on page 140.</p> |
| 10 | AWS RDS Oracle | <p>Consider using incremental backups for better backup performance.</p> <p>See “Protecting PaaS assets in AWS” on page 138 on page 138.</p> |
| 11 | AWS RDS Oracle | <p>Consider using archive redo log backups for better backup performance.</p> <p>See “About archive redo log backup for PaaS workloads” on page 140.</p> |

Table 2-1 Steps to protect PaaS assets in AWS (*continued*)

| Step no | Applicable to? | Steps |
|---------|----------------|--|
| 12 | All databases | 1 Discover your assets. See “Discovering PaaS assets” on page 141. 2 View your assets. See “Viewing PaaS assets” on page 142. |
| 13 | All databases | Manage credentials for your assets. See “Managing PaaS credentials” on page 142. |
| 14 | All databases | Add protection to your assets. See “Add protection to PaaS assets” on page 150. |
| 15 | All databases | Perform backup now, as required. See “Perform backup now” on page 150. |
| 16 | All databases | Recover PaaS assets. See “Recovering PaaS assets” on page 178. |
| 17 | All databases | View troubleshooting steps, if required. See “Troubleshoot PaaS workload protection and recovery issues” on page 216. |

Protecting PaaS assets in Azure

Follow these steps to protect your PaaS assets in Azure. The "Applicable to?" column shows for which database or environment a particular step is required. The Steps column gives you the links to the steps that you need to perform. Ensure that you perform each step completely before you proceed to the next step.

The steps that are applicable to all databases are mandatory steps. You need to perform the database-specific steps only when you protect that database.

Table 2-2 Steps to protect PaaS assets in Azure

| Step no | Applicable to? | Steps |
|---------|----------------------------------|---|
| 1 | All databases | Review the prerequisites for NetBackup PaaS protection. See “Prerequisites for protecting PaaS assets” on page 103. |
| | Azure Managed Instance databases | For Azure Managed Instance databases, consider these prerequisites and configure the required permissions. See “Protecting Azure Managed Instance databases” on page 110. |
| 2 | All databases | Consider the limitations and considerations of NetBackup protection for the databases that you want to protect. See “Limitation and considerations” on page 112. |
| 3 | All databases | Install the native client utilities. Each native client utility supports a set of databases. You can find the list of supported databases for each utility in the relevant sections. See “Installing the native client utilities” on page 129. |
| 4 | All databases | Configure storage for different NetBackup deployments. See “Configuring storage for different deployments” on page 136. |
| 5 | All databases | Configure the storage server for instant access. See “Configuring the storage server for instant access” on page 137. |

Table 2-2 Steps to protect PaaS assets in Azure (*continued*)

| Step no | Applicable to? | Steps |
|---------|---|--|
| 6 | All databases | NetBackup uses the NetBackup Snapshot Manager component to provide protection to the PaaS databases. Make sure that you do the following: <ul style="list-style-type: none"> ■ Install and configure NetBackup Snapshot Manager. ■ Install and configure the Azure plug-in for NetBackup Snapshot Manager. ■ Configure the required permissions for Azure in the NetBackup Snapshot Manager. See the <i>NetBackup™ Snapshot Manager for Cloud Install and Upgrade Guide</i> for more details. |
| 7 | When NetBackup is deployed in an AKS cluster. | Enable backup and restore in AKS. See “Enabling backup and restore in Kubernetes” on page 106. |
| 8 | MySQL and MariaDB databases | Enable binary logging. See “Enabling binary logging for MySQL and MariaDB databases” on page 106. |
| 9 | All databases | Consider Auto Image Replication (AIR) for extra protection. See “About Auto Image Replication for PaaS workloads” on page 140. |
| 10 | Azure SQL Server and Azure SQL Managed Instance | Consider using incremental backups for better backup performance. See “Protecting PaaS assets in AWS” on page 138 on page 138. |
| 11 | Azure MySQL server | Configure incremental backups for better backup performance. See “Configuring incremental backups for Azure MySQL server” on page 138. |

Table 2-2 Steps to protect PaaS assets in Azure (*continued*)

| Step no | Applicable to? | Steps |
|---------|----------------|--|
| 12 | All databases | 1 Discover your assets. See “Discovering PaaS assets” on page 141. 2 View your assets. See “Viewing PaaS assets” on page 142. |
| 13 | All databases | Manage credentials for your assets. See “Managing PaaS credentials” on page 142. |
| 14 | All databases | Add protection to your assets. See “Add protection to PaaS assets” on page 150. |
| 15 | All databases | Perform backup now, as required. See “Perform backup now” on page 150. |
| 16 | All databases | Recover PaaS assets. See “Recovering PaaS assets” on page 178. |
| 17 | All databases | View troubleshooting steps, if required. See “Troubleshoot PaaS workload protection and recovery issues” on page 216. |

Protecting PaaS assets in GCP

Follow these steps to protect your PaaS assets in GCP. The "Applicable to?" column shows for which database or environment a particular step is required. The Steps column gives you the links to the steps that you need to perform. Ensure that you perform each step completely before you proceed to the next step.

The steps that are applicable to all databases are mandatory steps. You need to perform the database-specific steps only when you protect that database.

Table 2-3 Steps to protect PaaS assets in GCP

| Step no | Applicable to? | Steps |
|---------|----------------|---|
| 1 | All databases | Review the prerequisites for NetBackup PaaS protection. See “Prerequisites for protecting PaaS assets” on page 103. |
| 2 | All databases | Consider the limitations and considerations of NetBackup protection for the databases that you want to protect. See “Limitation and considerations” on page 112. |
| 3 | All databases | Install the native client utilities. Each native client utility supports a set of databases. You can find the list of supported databases for each utility in the relevant sections. See “Installing the native client utilities” on page 129. |
| 4 | All databases | Configure storage for different NetBackup deployments. See “Configuring storage for different deployments” on page 136. |
| 5 | All databases | Configure the storage server for instant access. See “Configuring the storage server for instant access” on page 137. |

Table 2-3 Steps to protect PaaS assets in GCP (*continued*)

| Step no | Applicable to? | Steps |
|---------|---|--|
| 6 | All databases | <p>NetBackup uses the NetBackup Snapshot Manager component to provide protection to the PaaS databases. Make sure that you do the following:</p> <ul style="list-style-type: none">■ Install and configure NetBackup Snapshot Manager.■ Install and configure the GCP plug-in for NetBackup Snapshot Manager.■ Configure the required permissions for GCP in the NetBackup Snapshot Manager. <p>See the <i>NetBackup™ Snapshot Manager for Cloud Install and Upgrade Guide</i> for more details.</p> |
| 7 | When NetBackup is deployed in a Kubernetes cluster. | <p>Enable backup and restore in Kubernetes.</p> <p>See “Enabling backup and restore in Kubernetes” on page 106.</p> |
| 8 | MySQL and MariaDB databases | <p>Enable binary logging.</p> <p>See “Enabling binary logging for MySQL and MariaDB databases” on page 106.</p> |
| 9 | All databases | <p>Consider Auto Image Replication (AIR) for extra protection.</p> <p>See “About Auto Image Replication for PaaS workloads” on page 140.</p> |
| 10 | GCP SQL Server | <p>Consider using incremental backups for better backup performance.</p> <p>See “Protecting PaaS assets in AWS” on page 138 on page 138.</p> |

Table 2-3 Steps to protect PaaS assets in GCP (*continued*)

| Step no | Applicable to? | Steps |
|---------|----------------|--|
| 11 | All databases | 1 Discover your assets. See “Discovering PaaS assets” on page 141. 2 View your assets. See “Viewing PaaS assets” on page 142. |
| 12 | All databases | Manage credentials for your assets. See “Managing PaaS credentials” on page 142. |
| 13 | All databases | Add protection to your assets. See “Add protection to PaaS assets” on page 150. |
| 14 | All databases | Perform backup now, as required. See “Perform backup now” on page 150. |
| 15 | All databases | Recover PaaS assets. See “Recovering PaaS assets” on page 178. |
| 16 | All databases | View troubleshooting steps, if required. See “Troubleshoot PaaS workload protection and recovery issues” on page 216. |

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 2-4 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 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, DynamoDB, RDS Custom for Oracle, RDS Custom for SQL, AWS DocumentDB, and AWS Neptune. |
| GCP | SQL, PostgreSQL, MySQL, and BigQuery. |

Supported platforms

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

Table 2-5 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 AWS Neptune and the AWS RDS workloads support custom ports along with the default port.

Table 2-6 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 |
| AWS MariaDB | 3306 |
| AWS RDS AuroraDB Postgres | 5432 |
| AWS RDS AuroraDB MySQL | 3306 |
| AWS RDS SQL server | 1433 |
| AWS RDS for Oracle | 1521 |
| AWS DocumentDB | 27017 |
| AWS Neptune | 8182 |
| RDS Custom for Oracle | 1521 |
| RDS Custom for SQL | 1433 |

Table 2-6 Supported ports for PaaS (*continued*)

| Database PaaS workload | Supported ports |
|-----------------------------|-----------------|
| Azure Cosmos DB for NoSQL | 443 |
| Azure Cosmos DB for MongoDB | 10255 |
| GCP BigQuery | NA |
| GCP SQL Server port | 1433 |
| Amazon Redshift | 5439 |

Enabling binary logging for MySQL and MariaDB 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
```

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`

Protecting RDS Custom instances

NetBackup lets you protect RDS Custom for SQL Server and RDS Custom for Oracle databases, using the database native exports.

Protecting RDS Custom for SQL Server assets

To protect RDS Custom For SQL Server instances, deploy NetBackup on the EC2 instance in which RDS Custom SQL Server instances are running. NetBackup discovers these instances under the Microsoft SQL Server workload. You can perform instance-level and database-level backups using the MS-SQL-Server policy type and Protection plans. Refer to the *NetBackup for Microsoft SQL Server Administrator Guide* for more details regarding policy configuration, performing backup, and restore operations.

Consideration for protecting RDS Custom for SQL Server assets

- You can perform FULL, Differential incremental, and Transaction Log Backups of instances and databases.
- Restore is only supported at the individual database level.
- You can restore individual databases to their original location, alternate location, and alternate path. You can also restore the databases to an alternate RDS Custom SQL server.
- For restoring a database to an Amazon RDS Custom SQL Server, the restore destination path must be under the `D:\rdsdbdata` folder.

- Configurations for instant access are not supported while restoring databases to an Amazon RDS Custom SQL Server.

Protecting RDS Custom for Oracle assets

To protect RDS Custom for Oracle databases, deploy NetBackup on the EC2 instance where the RDS Custom Oracle is running. These instances are discovered under the Oracle workload. NetBackup uses RMAN to perform the backup and restore operation. See the *NetBackup for Oracle Administrator's Guide*, for details of policy configuration, backup, and restore operations.

Consideration for protecting RDS Custom for Oracle assets

- You can perform credential validation using the **OS authentication only** option, for the Oracle user.
- You must disable the **Archived redo logs in full and incremental schedules** option in the Oracle policy, to run full backups of RDS Custom Oracle databases.
- Full, Differential incremental backup, Cumulative incremental backup, and Archived redo log backup schedules are supported.
- To perform Archive log backup, set the **archivedLogRetentionHours** to maximum value to retain archive logs for a maximum period on the client RDS Custom Oracle instance.

Do the following:

- Create a text file named:
`/opt/aws/rdscustomagent/config/redo_logs_custom_configuration.json.`
- Add a JSON object in the following format:

```
{"archivedLogRetentionHours" : "num_of_hours"}
```

The number must be an integer in the range of 1–840.
- Locate the sample scripts for backup and restore, in the folder:
`/usr/opensv/netbackup/ext/db_ext/oracle/samples/rman/samples`
Copy the sample scripts to a different directory on your client. Modify script for your environment.
Update the **ORACLE_HOME**, **ORACLE_SID**, and **ORACLE_USER** fields in the script.
- Create a policy for the clients with the scripts option for performing cold and hot database backups using scripts.
- Backup and restore are supported only with RMAN restore script.
- Complete database restore and PIT database restore scripts are supported.

- Disable deletion protection of the RDS custom Oracle databases for successful backup and restore operations.
- For successful restore, move the RDS Custom Oracle instance in an automation paused state.

Protecting Azure Managed Instance databases

NetBackup version 11.0 onwards supports Azure Managed Instance backups using the native backup mechanism T-SQL. This backup method does not require creating a temporary database during backup. A `.bak` file is generated in the staging area on the Azure Blob storage and later moved to Universal share. The Managed Instance uses Azure SAS tokens to securely put the data in the form of `.bak` onto the Azure Blob. This method is supported only for full backups. If you use an AMI, the same AMI must be attached to the storage account. The backup and recovery workflows of these databases are the same as the other databases.

Prerequisites for protecting Azure Managed Instance databases

These prerequisites are required to protect Azure Managed Instance databases without creating a temporary database.

Note: Managed Instances use Azure SAS tokens to securely put the data in the form of a `.bak` file in the Azure Blob.

- Create a storage account name, using the following example:
 - Calculate the hash of the string formed by combining the Subscription ID and Region. Follow this example command:

```
echo -n "aaaaaaa-bbb-ccc-dd11-aabbccdd11223344eastus" |  
sha256sum  
2e5fb564552100a6060794937fedc33aefe4d01eb7e27971c1a0ed52c3f78b6c
```
 - Take only the first 16 chars of the hash, and prefix the characters **nbu**.
 - Your string becomes: **nbu + 2e5fb564552100a6**.
 - So, the Azure storage account name is: **nbu2e5fb564552100a6**
- For containers, use the following name format:
 - `azure-netbackup-<region>`
 - If the container does not exist, the NetBackup agent creates it depending on the availability of permissions.

- Management of encryption keys can be according to your organization.

Permissions required for protecting Azure Managed Instance databases

These permissions are required to protect Azure Managed Instance databases without creating a temporary database.

Note: Managed Instances use Azure SAS tokens to securely put data in the form of `.bak` files in the Azure Blob.

- Storage/storageAccounts/write
- Storage/storageAccounts/delete
- Storage/storageAccounts/listKeys/action
- Storage/storageAccounts/regenerateKey/action
- Storage/storageAccounts/read
- Storage/storageAccounts/blobServices/containers/read
- Storage/storageAccounts/blobServices/containers/write
- Storage/storageAccounts/blobServices/containers/delete
- Storage/storageAccounts/blobServices/containers/read
- Storage/storageAccounts/blobServices/generateUserDelegationKey/action
- Storage/storageAccounts/blobServices/containers/blobs/write
- Storage/storageAccounts/blobServices/containers/blobs/delete
- Storage/storageAccounts/blobServices/containers/blobs/read

Prerequisites for database level discovery

Follow these prerequisites to enable database level discovery for AWS RDS SQL, RDS PostgreSQL, and Aurora PostgreSQL.

- On the left, click **Cloud**, under **Workloads**. In the **Applications** tab, select and assign credentials to the assets that you want to protect. Click **Manage credentials** to assign credentials to the selected asset.
- After assigning credentials to the assets, select the asset and click **Discover**. The discovered databases contained in the discovered assets appear in the **PaaS** tab.

- In the **PaaS** tab, you can select individual databases that you want to protect.

Limitation and considerations

Consider the following when protecting PaaS workloads.

For all databases

- NetBackup Snapshot Manager deployment in RHEL 7.x is not supported for PaaS 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 184.
- 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.
- If full or differential incremental backups fail with temporary objects, then delete the temporary objects manually and run backup again.
- If the database has RLS (Row-Level Security), you must grant the backup user the permission to bypass all RLS policies, using the following command:

```
ALTER USER <backup_user> WITH BYPASSRLS;
```

For incremental backups for Azure PostgreSQL

- You can protect Azure PostgreSQL server assets with a protection plan and policy. You can use full and differential incremental schedules.
- Always set the `wal_level` server parameter as **logical**.
- The backup user must be an admin user.
- The backup user must have the `CREATEROLE` and `REPLICATION` permissions.
- To backup tables of other users, the tables must have the primary key.
- Intelligent groups are supported.
- Replica server databases are not supported for full and incremental backups.
- Large object data type is not supported for incremental backups.
- Incremental backup with the `bytea_output` server parameter with an escape value is not supported.
- Azure SMI and UMI are not supported for incremental backups.
- The last value of the sequence generator may not be consistent after restore of an incremental backup.
- To rename a database, do the following:
 - Note the name of the replication slot created for the database.

```
SELECT slot_name FROM pg_replication_slots WHERE database = <database name that needs to be renamed>
```
 - Drop the replication slot after renaming the database.

```
SELECT pg_drop_replication_slot('<replication slot name>')
```

For Amazon RDS PostgreSQL and Amazon 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://support.cohesity.com/s/article/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 Amazon 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 Amazon DocumentDB

- Only snapshot-based protection and restore from snapshots are supported.
- You can protect assets using NetBackup policies only. Protection plans are not supported.
- NetBackup replication feature is not supported.
- Intelligent groups are not supported.
- You can create IG for DocumentDB under cloud assets.
- NetBackup policies with the snapshot option do not list DocumentDB IG for backup selection.

For Amazon Neptune

- Only snapshot-based protection and restore from snapshots are supported.
- You can protect assets using NetBackup policies only. Protection plans are not supported.
- NetBackup replication feature is not supported.

- Intelligent groups are not supported.
- You can create IG for Neptune under cloud assets.
- NetBackup policies with the snapshot option do not list Neptune IG for backup selection.

For Amazon RDS SQL

- For credential validation, IAM is not supported for Amazon RDS SQL. You can use the username and password method.
- Cannot restore databases that contain a FILESTREAM file group.
- Cannot restore databases with the same name as an existing database. Database names must be unique.
- You can run up to two backups or restore tasks simultaneously for a specific RDS SQL instance.
- RDS SQL supports native restores of databases up to 16 TB. You can only restore 10 gigabytes of databases on SQL Server Express Edition.
- Native backups are not supported during the maintenance window or when Amazon RDS SQL is taking a snapshot of the database. A native backup task is cancelled, if it overlaps with the daily backup window of the RDS.
- The Transparent Data Encryption (TDE) certificate from the source RDS SQL instance must be present on the target RDS SQL instance to perform an alternate location restore of a TDE-enabled database.
- You can create native backups of TDE-enabled databases, but cannot restore those backups to on-premises databases.
- You can restore up to 10 user Transparent Data Encryption (TDE) certificates. These TDE certificates are used to restore existing backups. After restoring, these databases are encrypted using the master certificate. For more information about TDE certificates, see Amazon documentation.
- Do not start the TDE certificate names with RDSTDECertificate.
- You must add the SQL_BACKUP_AND_RESTORE option, to the option group attached to the RDS instance.

For Azure, Amazon RDS, and Aurora 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 cumulative incremental backups using Amazon RDS SQL

- You can protect Amazon RDS SQL database assets with a NetBackup protection plan or policy. You can use both full and cumulative incremental backup schedules.
- Backup and restore of TDE Certificates on RDS instances with Multi-AZ (Multi-Availability Zone) deployments is not supported directly on the RDS instance.
- To backup TDE certificates, disable Multi-AZ on the instance temporarily, and enable it after the backup is complete.
- To restore user TDE certificates, disable Multi-AZ on that instance temporarily, and enable it after the restore is complete.
- You cannot disable the automated backups when Multi-AZ is enabled on the RDS instance.
- Automated backups occur daily, during a user-configurable 30-minute window. These backups create more frequent requirements for full backups. A full backup runs at least once a day if the cumulative schedule comes after the automated backups.
- You can restore a full backup to Multi-AZ instance in the recovery mode only.
- Cumulative incremental backups can only be restored to an RDS instance in a single availability zone.
- Whenever a cleanup of backup history is performed either by calling the `sp_delete_backuphistory` API or by the Amazon maintenance window, the subsequent backup falls back to a full backup.
- When you perform a full backup of an Amazon RDS SQL Server asset using a different policy, the subsequent cumulative backup in the current policy is treated as full backups. This happens because the new full backup updates the `backupset ID` of the last full backup, causing the cumulative backup in the current policy to lose its reference point and run as a full backup instead.

For incremental backups using Amazon DynamoDB

- Use NetBackup policy to protect the Amazon DynamoDB table assets with incremental backups.
- You must enable point-in-time recovery in the Amazon Portal before configuring incremental backups. Enabling point-in-time recovery incurs an additional cost of approximately \$0.20 per GB-month. For example, if your table is 100GB, the monthly cost is approximately $100\text{GB} \times \$0.20 / \text{GB-month} = \20 .
- NetBackup enforces minimum and maximum incremental schedule intervals based on Amazon DynamoDB export limitations:
 - Minimum interval: 15 minutes. Incremental backups fail if the backup time window is less than 15 minutes.
 - Maximum interval: 24 hours per export operation. NetBackup performs multiple exports if the incremental backup time window is more than 24 hours.
- NetBackup performs a full back up during an incremental schedule, if the point-in-time recovery is disabled once and re-enabled for the table.

For incremental backups using Azure MySQL server

- You can protect Azure MySQL server assets with a protection plan and policy. You can use full and differential-incremental schedules at the instance level. The individual databases can only be protected using the full schedule.
- You can recover individual databases from the backups of a server taken to another destination server. Restore fails if a database with the same name exists on the destination server.
- Currently, NetBackup does not restore the users and their permissions to the destination server. All database objects are recovered with the same users as in the source database during the time of the backup. You can create users and grant the necessary permissions, after restore.
- Records in the tables created with storage engine type MEMORY are not backed up during an incremental schedule. These records remain in memory, and changes made to those tables do not reflect in the binary logs.
- NetBackup performs a full backup during an incremental schedule in the following scenarios:
 - One or more binary logs on the server are purged between the subsequent incremental backups. Ensure that the `binlog_expire_logs_seconds` value is set to the appropriate value based on the frequency of the incremental schedule.

- If you change the schema of one or more databases on a server and then perform any DDL action on any of those databases.
- One or more databases are added or removed from the server.
- If the server is configured as High Availability, and a failover occurs on the server.
- If the asset's maximum incremental recovery points criteria (100) is reached under the subscribed policy or protection plan.

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 Incremental backups using GCP PostgreSQL

- You can protect GCP Postgres SQL assets with NetBackup protection plan or policy.
- Both full and differential incremental schedules are supported.
- To enable logical decoding on Google Cloud SQL for PostgreSQL, configure the database flag `cloudsql.logical_decoding`. Do the following:
 - Navigate to your Cloud SQL instance in the GCP console.
 - Click Edit, and scroll to the Flags section.
 - Add or modify the parameter `cloudsql.logical_decoding` flag.
 - Save the changes and restart the instance. You must restart of the Cloud SQL instance for these changes to take effect.
- Assign REPLICATION permissions to the backup user.
- Replica server databases are supported for full and incremental backup.
- Databases having Large objects are not supported for incremental backup.
- Incremental backups are not supported with the `bytea_output` server parameter having the value: `escape`.
- Incremental backups are not supported, if:
 - If the backup user is non-admin.
 - The tables of other users do not have a primary key.
- After restoring an incremental backup, the last value of the sequence may not be consistent.
- Before renaming a database, follow these steps:
 - Note down the name OS the replication slot created for the database that you want to rename.
 - Select the `slot_name` form `pg_replication_slots`, where: `database = <database name that needs to be renamed>`
 - Drop the replication slot after renaming the database: `SELECT pg_drop_replication_slot(<replication slot name>)`

For Incremental backups with Amazon RDS PostgreSQL and Amazon Aurora PostgreSQL

- You can protect Amazon RDS PostgreSQL and Amazon Aurora PostgreSQL assets with NetBackup protection plan or policy.
- Both full and differential incremental schedules are supported.
- To enable logical decoding in Amazon RDS PostgreSQL and Aurora PostgreSQL, configure the `rds.logical_replication` server parameter to `1`, in the custom DB parameter group.
Do the following in the AWS console:
 - Create a custom DB parameter group (for RDS) or custom DB cluster parameter group (for Aurora).
 - Set the `rds.logical_replication` server parameter to `1`.
 - Associate the parameter group with your database instance or cluster.
 - Restart the writer instance to apply changes.
 - Verify settings using SQL: `SHOW wal_level`.
- Replica server databases are not supported for full and incremental backup.
- Databases having Large objects are not supported for incremental backup.
- Incremental backups are not supported with the `bytea_output` server parameter having the value: `escape`.
- Incremental backups are not supported, if:
 - If the backup user is non-admin.
 - The tables of other users do not have a primary key.
- After restoring an incremental backup, the last value of the sequence may not be consistent.
- Before renaming a database, follow these steps:
 - Note down the name OS the replication slot created for the database that you want to rename.
 - Select the `slot_name` form `pg_replication_slots`, where: `database = <database name that needs to be renamed>`
 - Drop the replication slot after renaming the database: `SELECT pg_drop_replication_slot(<replication slot name>)`

For Azure SQL and SQL Managed Instance

These limitations and considerations apply to Azure SQL database and Azure Managed Instance backups.

- **Configure the backup method**— you can create and configure the `dbpaas_config.json` file for each policy to specify a backup method to fit your requirements. Note that you must give the CDC consent through NetBackup UI.

- **Create the `dbpaas_config.json` file at:**

```
/openv/netbackup/db/>dbpaas_config.json
```

Follow an equivalent path for Windows OS.

- **Here is a sample `dbpaas_config.json` file:**

```
{
  "Configurations": [
    {
      "PolicyName": "policy_name",
      "BackupMethod": "OFFLOAD_TO_BACKUP_HOST",
      "BatchCount": 1024
    }
  ]
}
```

- **Explanation of the parameters:**

- `PolicyName`—the name of the policy to which you want to apply the JSON.
- For `BackupMethod` use any of the following values:
 - `OFFLOAD_TO_SERVER`—this performs all the CDC processing at the SQL/MI server side. This is the default behavior.
 - `OFFLOAD_TO_BACKUP_HOST`—in this method, the CDC data is fetched to the backup host from the server. The fetched data is processed at the backup host.
- `BatchCount`—determines how many transactions of records are processed at a time.
- For Azure SQL database and Azure SQL Managed Instance workloads, NetBackup protection plans do not include some backup configuration options such as `BackupMethod` and `BatchCount`. To use these options:
 - Use a NetBackup policy instead of a protection plan.

- Reference the configured policy in your DBPaaS configuration file.
When using a protection plan, any changes to these options in the DBPaaS configuration file are ignored.
- If an Azure VM is used as a media server, it must 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.
- NetBackup requires user consent to enable SQL Server's Change Data Capture (CDC) for backup operations. This CDC option is enabled by default in the protection plans.
 - If the CDC option is enabled, NetBackup uses SQL Server's CDC mechanism.
 - If the CDC option is disabled, backups fail—except for Azure SQL Managed Instances using Customer-Managed Keys (CMK) and full backup schedules.
- After deleting a database, you must stop CDC capture, disable cleanup jobs, and perform manual cleanup.
- For smooth CDC performance, do the following:
 - Do not use DDL triggers that block schema changes. Like, *CREATE/DROP TABLE*, *CREATE/DROP FUNCTION*.
 - Ensure that the SQL Server Agent service is running.
 - Use an account with db_owner or sysadmin privileges for CDC setup.
 - To avoid capture_instance conflicts, check if the table is already CDC-enabled.
- Backup fails when Readlock is placed on the database or resource group.
- If the databases contain any of these types of tables, the backup fail due to CDC limitations.
 - Graph tables
 - Temporal tables
 - Ledger tables (Updatable ledger)
 - Memory optimized table (business critical tier only)
For Azure Managed instance databases, if the backup taken using the native backup database workflow, using TDE enabled by Customer-managed Keys or TDE disabled, these types of tables are supported.
- Database diagrams not restored.

- NetBackup can back up and restore databases to new environments, but the identity columns may not retain their original seed values. This can cause new records to receive unexpected identity numbers, especially if rows were deleted before the backup. However, this issue does not affect TDE-enabled Azure SQL Managed Instance databases configured with a CMK and backed up either with the native BAK format or through a policy without CDC consent.
- NetBackup supports column-level encryption only for Azure SQL Managed database instances where TDE is configured with CMK. Backups must be full backups, taken either with the native BAK format or through a policy without CDC consent.
- 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.
- Consider the following for the SqlPackage utility.
 - SqlPackage may fail to export the following SQL Server objects: Logins, Jobs (SQL Agent Jobs), Linked Servers, Server-level permissions, Server roles, Credentials, Database Mail configuration, Replication settings, SSIS packages, Service Broker objects, CLR assemblies, Data in FILESTREAM/FILETABLE, Encrypted objects (if keys/certs are not exported), Cross-database dependencies, Extended stored procedures, and certain system objects.
 - SqlPackage primarily exports database-level objects and data, but not server-level or agent-level configurations. For a complete migration, use additional tools or scripts to handle these objects.

For Azure SQL Server and SQL Managed Instance incremental backup

- Set the retention period of CDC cleanup jobs to match the incremental schedule and keep them until the next full backup. Otherwise, subsequent backups revert to full.
- 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 higher. Incremental backups of Azure SQL Managed Instances can run only on NetBackup media server version 10.3 and above.

- 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.
- 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 Amazon 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.
- EFS support as a staging location is discontinued from NetBackup 11.0.0.1 release. Use S3 as a staging location for RDS Oracle backups. You must manually clean up the EFS mounts that were previously used as staging locations. Any staging cleanup job run from the NetBackup UI, only removes the location from the NetBackup database.
- Before running the archive log backup, set the retention period in the protection plan. See the Knowledge-base article: <https://support.cohesity.com/s/article/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 databases

- Redshift database restores to alternate regions or alternate accounts are not supported.
- External databases created through Glue Registration are not supported for backup and recovery, even though they may appear in the PaaS assets list.
- FIPS is currently not supported for Redshift databases.
- 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.
- Redshift cluster must be in available state before starting a database backup.
- 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.

For Amazon Redshift clusters

- The minimum supported version of Primary, Media, and Snapshot management server is NetBackup 10.5.
- FIPS is currently not supported for Redshift clusters.
- Redshift clusters created with AWS Secrets Manager credentials are not supported.
- Redshift cluster restores to alternate regions or alternate accounts are not supported.
- Redshift cluster must be in available state before triggering a cluster snapshot.
- Cluster restore jobs may appear successful immediately in the Activity monitor, while the job is still in progress. Monitor the cluster restore job in your AWS console for the actual status of the job.
- Maximum number of manual snapshots per Redshift cluster is 20.

- During restore, the `PubliclyAccessible` property of the restored cluster is set to `False`. You can manually change it as required after the restore.
- Do not trigger NetBackup image expiry for Redshift cluster snapshot images when a restore is in progress. If automated image expiry jobs run while a restore is ongoing, then the cleanup of the snapshot from the AWS portal fails.
- NetBackup Activity monitor does not show values for these snapshot-parameters: Bytes transferred, Files written, Current File, Estimated Files Remaining, and Estimated Files.

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 GCP BigQuery

- GCP Project users and their permissions are not restored. The restored dataset owner is the GCP service principal configured in NetBackup, while adding the cloud provider.
- The maximum data export limit is 50 TB per day, per project.
- Tags attached to datasets and tables are not restored.
- The datasets created for multiple regions (US, EU) are not supported.
- The RANGE data type export is not supported by the GCP export API.
- The linked datasets are not discovered, backed up, or restored.
- A record with the data type: DATETIME in the tables is not supported for backup.
- NetBackup does not back up tables without any column or schema.
- Backups of tables with column-level access control or row-level security are not supported.

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

You need to install this utility to protect the following databases:

- Azure MySQL
- Azure MariaDB
- AWS MySQL
- AWS MariaDB
- AWS Aurora MySQL
- GCP MySQL

The utility version must be the same or higher than the server version.

- For server version 8.0, the recommended version of the MySQL client utility is 8.0.42.
- For server version 8.4, the recommended version of the MySQL client utility is 8.4.4.

RPM download locations:

Server version 8.4 https://downloads.mysql.com/archives/get/p/23/file/mysql-community-client-8.4.4-1.el9.x86_64.rpm

Server version 8.0 https://downloads.mysql.com/archives/get/p/23/file/mysql-community-client-8.0.42-1.el9.x86_64.rpm

To install, do the following in the terminal:

- 1 Navigate to `/usr/opencv/dbpaas`. If the `dbpaas` folder does not exist, run the command `mkdir dbpaas` to create the folder.
- 2 Create the `mysql` folder inside the `/usr/opencv/dbpaas` folder. Run the command `mkdir mysql` inside `/usr/opencv/dbpaas`.
 - For server version 8.0, create the folder `80` inside the `/usr/opencv/dbpaas/mysql` folder. Run the command `mkdir 80` inside the `/usr/opencv/dbpaas/mysql` folder.
 - For server version 8.4, create the `84` folder inside the `/usr/opencv/dbpaas/mysql` folder. Run the command `mkdir 84` inside the `/usr/opencv/dbpaas/mysql` folder.
- 3 Download the utilities as follows:
 - For server version 8.0, go to the `80` folder, and run the command:

```
wget https://downloads.mysql.com/archives/get/p/23/file/mysql-community-client-8.0.42-1.el9.x86_64.rpm
```
 - For server version 8.4, go to the `84` folder, and run the command:

```
wget https://downloads.mysql.com/archives/get/p/23/file/mysql-community-client-8.4.4-1.el9.x86_64.rpm
```
- 4 To install, run the command:

For version 8.0:

```
rpm2cpio
/usr/opencv/dbpaas/mysql/80/mysql-community-client-8.0.42-1.el9.x86_64.rpm
| bsdtar -xv --strip-components=2 -f - -C
/usr/opencv/dbpaas/mysql/80 'usr/bin/mysql' 'usr/bin/mysqldump'
'usr/bin/mysqlbinlog'
```

For version 8.4:

```
rpm2cpio
/usr/opencv/dbpaas/mysql/84/mysql-community-client-8.4.4-1.el9.x86_64.rpm
| bsdtar -xv --strip-components=2 -f - -C
/usr/opencv/dbpaas/mysql/84 'usr/bin/mysql' 'usr/bin/mysqldump'
'usr/bin/mysqlbinlog'
```

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

Installing the SqlPackage client utility

You need to install this utility to protect the following databases:

- Azure SQL
- Azure SQL Managed Instance
- AWS RDS SQL server
- GCP SQL server

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
```

If you still get the error: `sqlpackage` client utility does not exist on the given host, create a soft link for the `sqlpackage` to the path `/usr/bin/sqlpackage`.

For example:

If the `sqlpackage` is located at `/root/sqlpackage/sqlpackage`, create a soft link and then run the backup as:

```
ln -s /root/sqlpackage/sqlpackage/usr/bin/sqlpackage
```

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

You need to install this utility to protect the following databases:

- Azure PostgreSQL Single and Flexible server
- AWS RDS PostgreSQL
- AWS RDS Aurora PostgreSQL
- GCP PostgreSQL

The utility version must be the same or higher than the server version.

- For server version 16, the recommended version of the PostgreSQL client utility is 16.9.
- For server version 17, the recommended version of the PostgreSQL client utility is 17.5.

| | | |
|-------------------|-----|---|
| Server version 16 | Bin | https://download.postgresql.org/pub/repos/yum/16/redhat/rhel-9-x86_64/postgresql16-16.9-3PGDG.rhel9.x86_64.rpm |
| | Lib | https://download.postgresql.org/pub/repos/yum/16/redhat/rhel-9-x86_64/postgresql16-libs-16.9-3PGDG.rhel9.x86_64.rpm |
| Server version 17 | Bin | https://download.postgresql.org/pub/repos/yum/17/redhat/rhel-9-x86_64/postgresql17-17.5-3PGDG.rhel9.x86_64.rpm |
| | Lib | https://download.postgresql.org/pub/repos/yum/17/redhat/rhel-9-x86_64/postgresql17-libs-17.5-3PGDG.rhel9.x86_64.rpm |

To install, do the following:

- 1 Navigate to `/usr/opensv/dbpaas`. If the `dbpaas` folder does not exist, run the command `mkdir dbpaas` to create the folder.
- 2 Create the `postgresql` folder inside the `/usr/opensv/dbpaas` folder. Run the command `mkdir postgresql` inside `/usr/opensv/dbpaas`.
 - For server version 16, create the folder `16` inside the `/usr/opensv/dbpaas/postgresql` folder. Run the command `mkdir 16` inside the `/usr/opensv/dbpaas/postgresql` folder.

- For server version 17, create the folder 17 inside the `/usr/opensv/dbpaas/postgresql` folder. Run the command `mkdir 17` inside the `/usr/opensv/dbpaas/postgresql` folder.

3 Download the utilities as follows:

- For server version 16, go to the 16 folder, and run the command:

For bin

```
wget https://download.postgresql.org/pub/repos/
yum/16/redhat/rhel-9-x86_64/postgresql16-16.9-
3PGDG.rhel9.x86_64.rpm
```

For lib

```
wget https://download.postgresql.org/pub/repos/
yum/16/redhat/rhel-9-x86_64/postgresql16-libs-16.9
-3PGDG.rhel9.x86_64.rpm
```

- For server version 17, go to the 17 folder, and run the command:

For bin

```
wget https://download.postgresql.org/pub /repos/
yum/17/redhat/rhel-9-x86_64/postgresql17-17.5-
3PGDG.rhel9.x86_64.rpm
```

For lib

```
wget https://download.postgresql.org/pub/repos/  
yum/17/redhat/rhel-9-x86_64/postgresql17-libs-  
17.5-3PGDG.rhel9.x86_64.rpm
```

4 To install, run the commands:

For version 16 Lib:

```
rpm2cpio /usr/opensv/dbpaas/postgresql/16/postgresql16-  
libs-16.9-3PGDG.rhel9.x86_64.rpm | bsdtar -xv --strip-components=3  
-f - -C /usr/opensv/dbpaas/postgresql/16  
'usr/pgsql-16/lib/libpq.so*'
```

For version 16 Bin:

```
rpm2cpio  
/usr/opensv/dbpaas/postgresql/16/postgresql16-16.9-3PGDG.rhel9.x86_64.rpm  
| bsdtar -xv --strip-components=3 -f - -C  
/usr/opensv/dbpaas/postgresql/16 'usr/pgsql-16/bin/pg_dump'  
'usr/pgsql-16/bin/pg_dumpall' 'usr/pgsql-16/bin/psql'  
'usr/pgsql-16/bin/pg_restore' 'usr/pgsql-16/bin/createdb'  
'usr/pgsql-16/bin/dropdb'
```

For version 17 Lib:

```
rpm2cpio  
/usr/opensv/dbpaas/postgresql/17/postgresql17-libs-17.5-3PGDG.rhel9.x86_64.rpm  
| bsdtar -xv --strip-components=3 -f - -C  
/usr/opensv/dbpaas/postgresql/17 'usr/pgsql-17/lib/libpq.so*'
```

For version 17 Bin:

```
rpm2cpio  
/usr/opensv/dbpaas/postgresql/17/postgresql17-17.5-3PGDG.rhel9.x86_64.rpm  
| bsdtar -xv --strip-components=3 -f - -C  
/usr/opensv/dbpaas/postgresql/17 'usr/pgsql-17/bin/pg_dump'  
'usr/pgsql-17/bin/pg_dumpall' 'usr/pgsql-17/bin/psql'  
'usr/pgsql-17/bin/pg_restore' 'usr/pgsql-17/bin/createdb'  
'usr/pgsql-17/bin/dropdb'
```

Installing MongoDB client utility

You need to install this utility to protect Azure Cosmos DB for MongoDB databases.

The MongoDB client utility recommended version is 100.10.0.

Download <https://www.mongodb.com/try/download/database-tools/releases/archive>
location

To install, run the following command in the terminal:

```
rpm -ivh mongodb-database-tools-rhel70-x86_64-100.9.4.rpm
```

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/opencv/netbackup/vpfs_override_parameters.json
```

2. Update the `CloudCacheSize` parameter with the available storage size in GBs.

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

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 `polycoreutils` and `polycoreutils-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
```

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.

Configuring incremental backups for Azure MySQL server

NetBackup supports both full and differential incremental backup schedules for Azure MySQL server. During full backup, NetBackup takes a logical dump of all user databases. During incremental backups, NetBackup downloads the binary logs generated on the MySQL server between subsequent incremental backups. When you restore a database, NetBackup creates the database with the full backups and then applies the binary logs in chronological sequence.

Configuring server parameter

Configure the following parameters in the Azure portal.

Table 2-7 Parameters for incremental backup

| Parameter | Description | Recommendation |
|--|---|--|
| <code>binlog_expire_logs_seconds</code> | Number of seconds to wait before the binary log file is purged. | This value must be higher than the frequency of backups in the schedule. Ensures that the binary logs on the server are not purged between subsequent incremental backups. |
| <code>log_bin</code> | The status of binary logging on the server, either enabled (ON) or disabled (OFF). | For incremental backups, this parameter must be set to ON. It captures the changes in the binary log file. |
| <code>log_bin_trust_function_creators</code> | This variable applies when binary logging is enabled. It allows the stored function creators to be trusted to make stored functions that allow unsafe events to be written to the binary log. | To restore databases to servers, set this parameter as ON. |
| <code>default_table_encryption</code> | Specifies the default encryption setting to use for the schemas and general tablespaces, when they are created without specifying an ENCRYPTION clause. | To restore databases to servers, set this parameter as ON. |
| <code>max_binlog_cache_size</code> | If a transaction requires more memory than this specified value, the server generates a storage error. | Ensure that this value can accommodate the maximum total size of a transaction. We recommend keeping the default value of this parameter. |

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.

About Auto Image Replication for PaaS workloads

NetBackup supports Auto Image Replication (AIR) to a target LSU hosted locally or in the cloud. The LSU must have a standard or archive storage class or tier on the target domain with a one-to-many AIR model.

Currently the Amazon Redshift workloads are not supported for AIR.

You can configure AIR using a protection plan or a policy. Both support one-to-one replication.

One-to-many replication is supported only when AIR is configured using policy, as policy allows multiple replication targets.

Protection plans support only a single replication target.

While adding a schedule to a protection plan, select the **Replicate this backup** option. If AIR is enabled for a schedule in the protection plan, it is recommended to enable it for all schedules in that protection plan.

Note: For archive tier storage class LSU, only full backup schedule is supported.

You can also configure AIR using policy by selecting Storage Lifecycle Policy as the storage.

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 Web UI 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 Cohesity Technical Support.

To discover PaaS assets:

- 1 Add a Snapshot Manager. See [“Add a Snapshot Manager”](#) on page 15.
- 2 Add Microsoft Azure, GCP, or AWS as a provider. See [“Add a cloud provider for a Snapshot Manager”](#) on page 15.
- 3 Run a discovery. See [“Discover assets on Snapshot Manager”](#) on page 22.

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 GCP BigQuery, DynamoDB, Amazon Redshift, RDS Custom for Oracle, RDS Custom

for SQL, AWS DocumentDB, and AWS Neptune 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 145.

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 146.

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*.

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 > Microsoft Entra ID > 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.

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, SHOW_ROUTINE, EVENT, TRIGGER ON *.* TO `dbuser`@`%` WITH GRANT OPTION;`
- For Azure MySQL incremental protection add these additional permission:
`GRANT SET_USER_ID, REPLICATION CLIENT, SESSION_VARIABLES_ADMIN, REPLICATION_APPLIER 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.

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.

For RDS Custom for Oracle, RDS Custom for SQL, AWS DocumentDB, and AWS Neptune assets the add protection option is not available.

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 143..
If the **Credential name** column is empty, you need to assign a credential to the asset.
See [“Add credentials to a database”](#) on page 143.
- 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**.

The Redshift cluster, AWS DocumentDB, and AWS Neptune assets are not protected using a protection plan. You can protect them using policies. See [“Managing policies for cloud assets”](#) on page 39.

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.

Note the following:

- 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.

- For archive redo log type schedules, NetBackup takes full backup for backup now, irrespective of what you specified in the protection plan.
- For Redshift cluster, AWS DocumentDB, and AWS Neptune assets the backup now option is not available. You can use the Manual backup option in the policy, to initiate a backup.

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).

Recovering cloud assets

This chapter includes the following topics:

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

Recovering cloud assets

You can restore AWS, Azure, Azure Stack Hub, OCI, and GCP VM assets from snapshot copy, backup copy, or duplicate copy. For AWS, you can restore from a replica copy too. 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.
- (Azure and Azure Stack Hub) 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 Hub)
- 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 3-1 Supported parameters for Azure, Azure Stack Hub, GCP, OCI, and AWS snapshot and backup copies

| Parameters | Snapshot copy | | | | Backup copy | | | |
|---------------------------------------|---------------|-----------------|-------------|-----|-------------|-----------------|-------------|-----|
| | Azure | Azure Stack Hub | GCP and AWS | OCI | Azure | Azure Stack Hub | 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 | |

Table 3-1 Supported parameters for Azure, Azure Stack Hub, GCP, OCI, and AWS snapshot and backup copies (*continued*)

| | | | | | | | | |
|--|---|---|---|---|---|---|---|---|
| Change resource group | Y | Y | | | Y | Y | | |
| Change the region of the VM | | | | | Y | Y | Y | |
| 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 |

Restoring to a different cloud provider

NetBackup 10.5 or later, lets you restore protected cloud VMs in one cloud provider to another cloud provider, from the backup from snapshot copy. For example, AWS to Azure and vice versa.

To restore across providers, at least one media server must be 10.5 or higher. This media server must have access the STU where the backed-up image resides.

Follow these steps to restore VMs across different providers:

You need to perform the following steps to restore to a different cloud provider.

- Prepare the VMs that you want to back up.
- Perform back up of the VMs.
- Restore to the desired cloud provider.
- Perform the post-restore configurations.

For a detailed description of each of these steps, refer to the following section:

See [“Restore to a different cloud provider”](#) on page 163.

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 Hub 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 Hub 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.

Note: For ADE encrypted disks, the **Edit the encryption key** option would be disabled.

The **Azure disk encryption** column displays the status of ADE encryption.

- Select the required **Encryption type**.
- Select the required encryption **Key** and click **Save**.

8 Click **Next**.

9 In the Recovery options page:

- *(Only for Azure and AWS)* Select **Restore network configuration** option to restore the VM with the same network configuration as the source VM.
- To change the network configuration, select **Change network configuration** option, 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.

When VM is in **updating** provisioning state, the recovery job would not fail but would be in waiting state for 5 minutes for the status to change from updating to succeeded.

Note: For ADE enabled VM's, if VM is in **updating** provisioning state and extension is not installed, then VM creation would fail and resources would be cleaned up.

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

<https://support.cohesity.com/s/article/article-100040135>

Parallel restore for cloud VMs

NetBackup 11.2 introduces parallel restore for cloud VM workloads to reduce restore time for virtual machines with multiple disks. Parallel Restore improves the performance of cloud virtual machine recovery by restoring multiple disks in parallel instead of sequentially.

Supported platforms: AWS and Microsoft Azure

Prerequisites

Parallel restore is supported with NetBackup 11.2 or later, subject to the following version requirements:

- The datamover (NetBackup Snapshot Manager/DM) must be running version 11.2 or later.
- Primary and media servers must be version 11.2 or later.

Note the following:

Parallel restore is supported only when the primary server, media server, and datamover (DM) are all running NetBackup 11.2 or later.

- If the primary server is earlier than 11.2, NetBackup performs a sequential restore; parallel restore is not triggered, but the restore completes successfully.
- If the primary server is 11.2 or later and the media server is earlier than 11.2, the restore fails. In this case, the media server must be upgraded.
- When the primary server, media server, and datamover are all running 11.2 or later, NetBackup performs a parallel restore.

Parallel disk restore workflow

During a parallel restore operation, NetBackup coordinates the restore using a parent job and multiple child jobs. A single parent job controls the overall restore workflow and manages the pre-restore and post-restore phases. As part of the pre-restore phase, the parent job creates the required staging disks for the restore operation.

NetBackup then starts one restore stream per disk. Each restore stream identifies the required data extents for its assigned disk and restores the disk data independently. These restore streams run concurrently, allowing all disks associated with the virtual machine to be restored in parallel. The parent job monitors the progress of the child restore jobs and completes the restore workflow after all disk restores finish successfully and required resources are available.

Parallel restore automatically uses the available media server and storage capabilities to determine, at that time, the number of concurrent restore streams that can run and the number of streams that can be queued. No manual configuration is required to enable this behavior for supported cloud virtual machine workloads.

Parallel restore reduces overall recovery time and improves recovery time objectives (RTOs) for critical workloads.

By default, NetBackup enables parallel restore for supported cloud VM workloads.

Note: To disable parallel restore add the following parameter to `bp.conf` file:

```
CLOUDVM_PARALLEL_RESTORE_DISABLE = 1
```

Activity monitor will display the parent jobs and child jobs for each stream/disk.

Performance considerations

Parallel Restore improves restore performance, particularly for virtual machines with multiple data disks, based on available resources and storage capabilities.

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 the following:

- *(Applicable for Azure cloud)* Application restore to an alternate location for the ADE-enabled VMs is not supported.
- *(For OCI)* If a volume does not have a device name when the snapshot is taken, then, the original location restores of that volume is attached to the next available device on the original VM.

Additional steps required after restoring an AWS RDS database instance

The following steps are required after you restore an AWS RDS database instance snapshot. Even though the restore operation is successful, these manual steps are required so that the instance is available for normal use.

After restoring an AWS RDS database instance successfully, you have to manually check and reassign certain properties of the restored instance. This is required because even though the restore operation itself is successful, one or more instance properties are not restored completely. In some cases, NetBackup resets the property values to their default settings.

The following RDS database instance or cluster properties are not restored completely and needs modification:

- **VPC security groups** value (*AWS Management Console > RDS Database instance > Connectivity & security tab*)

- **Deletion protection** setting (*AWS Management Console > RDS Database instance > Configuration tab*)
- **Copy tags to snapshots** setting (*AWS Management Console > RDS Database instance > Maintenance & backups tab*)

Perform the following steps:

- 1 Verify that the RDS database instance snapshot restore is successful.
- 2 Sign in to the AWS Management Console and from the top-right corner, select the region in which you have restored the RDS instance.
- 3 From the Services menu, under Database, click **RDS**.
- 4 From the Dashboard menu on the left, click **Databases**.
- 5 In the Databases panel, select the restored RDS database instance and then click **Modify** from the menu bar on the top right.
- 6 On the Modify DB panel, check for the following properties and ensure that the attribute values match with those of the original instance:
 - Under Network & Security, verify that the **Security group** attribute has the correct security group name assigned.
 - Under Backup, verify that the **Copy tags to snapshots** option is set as per the original instance.
 - Under Deletion protection, verify that the **Enable deletion protection** option is set as per the original instance.
 - If required, verify all the other parameter values and set them as per your preference.
- 7 Once you have modified the desired RDS instance properties, click **Continue**.
- 8 Under Scheduling of modifications, choose an appropriate option depending on when you want to apply the modifications to the instance and then click **Modify DB instance**.
- 9 Verify the RDS instance properties and ensure that the changes have taken effect.

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 3-2 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\n MissingDependencies</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> |

(GCP only) Restoring virtual machines and volumes using the autoDelete disk support

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 the following:

- Snapshot replicas do not support rollback.
- Azure Stack Hub, 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.

Restore to a different cloud provider

NetBackup 10.5 or higher, lets you restore protected VMs from one cloud provider to another cloud provider from the backup from snapshot copy. For example, AWS to Azure and vice versa.

To restore across providers, at least one media server must be 10.5 or higher. This media server must have access the STU where the backed-up image resides.

[Table 3-3](#) provides the steps to be performed for the following operating systems:

- RHEL
- SLES
- Ubuntu
- Windows

Table 3-3 Steps to be performed to restore between AWS and Azure cloud

| Steps | Description |
|---|---|
| Step 1: Prepare the VMs for back up | <ul style="list-style-type: none"> ■ Check if the VM is part of supported VM (OS/Platform). Refer to the Hardware and Cloud Storage Compatibility List (HCL). ■ Back up the source VM before making any changes. ■ Identify and install the appropriate drivers that must be preinstalled before backing up the source VM. <p>Depending on the operating system, refer to the following specific sections for preparing the VMs:</p> <ul style="list-style-type: none"> ■ RHEL ■ SLES ■ Ubuntu ■ Windows |
| Step 2: Backup the VMs | Back up the VM after installing the drivers. |

Table 3-3 Steps to be performed to restore between AWS and Azure cloud *(continued)*

| Steps | Description |
|---|--|
| Step 3: Cross cloud provider restore configuration | <ul style="list-style-type: none"> ■ Select cross-provider configuration where you want to restore the target VM in the destination cloud provider. ■ Select the appropriate instance size of the target VM based on your requirement. <p>For target AWS:</p> <ul style="list-style-type: none"> ■ Ensure that you have a compatible AMI ID while restoring the VM. ■ Required parameters to restore to AWS: ImageId (AMI) and Instance type ■ Select AMI and instance type that supports the source VM's boot mode (Legacy BIOS/UEFI) For more information, refer to Amazon EC2 boot modes. ■ Matching OS platform with source VM is required. ■ Log on using same username and password as source VM. <p>Note: If the instance reachability check fails, then perform the steps mentioned in the following section after the VM is restored:</p> <p style="text-align: center;">Step 4: Post restore configuration</p> <p>For target Azure:</p> <p>Select instance type that supports the source VM's boot mode and supported number of attached disk drives.</p> |
| Step 4: Post-restore configurations | <p>Additional configuration to ensure that the restored instance is reachable.</p> <p>Depending on the operating system, refer to the following sections for post restore configurations:</p> <ul style="list-style-type: none"> ■ RHEL ■ SLES ■ Ubuntu ■ Windows |

Prepare the VMs for back up

This section describes the considerations and prerequisites to back up the VMs for restoring to a different cloud platform. The process is different for different operating systems, depending on the cloud service to which you want to restore.

RHEL

Target: AWS

1 Install the required Xen and Nitro drivers:

- If drivers are not installed, then install them by running the following commands:
 - `lsinitrd | grep -i -e nvme -e ena -e xen`
 - `modinfo nvme`
 For more information, refer to [Install or upgrade the NVMe driver](#)
 - `modinfo ena`
 For more information, refer to [Enable enhanced networking with the Elastic Network Adapter \(ENA\) on your EC2 instances](#)
 - Update/Create the `/etc/dracut.conf` file with the following line:
`add_drivers+="xen-blkfront xen-netfront nvme-core nvme"`
 - Run the following command:
`dracut -f -v`
- Validate if the drivers are installed successfully by running the following command:
`lsinitrd | grep -i -e nvme -e ena -e xen`

2 To avoid mount failures, it is recommended to replace the device names with **UUID** in `/etc/fstab` file.

Back up the original `fstab` file and comment out Azure-specific entries, and other non-critical entries which might cause boot failures after restore. Alternatively, you can add `nofail` in `fstab` file for these entries.

3 Create a root user password.

4 Configure or obtain the root user credentials, if the VM is configured with key based logon.

To use the key based logon, perform the following:

- Back up the original `/root/.ssh/authorized_keys`.
- The `/root/.ssh/authorized_keys` contains the same public key as **azureuser**, but cannot log on using the root user and key due to the following command present in `authorized_keys` for the root user and the associated key:

```
`echo 'Please login as the user \"azureuser\" rather than the user \"root\".';echo;sleep 10;exit 142`
```

Note: This is applicable for Azure-created keys and user-provided keys.

You must delete the command to let the root logon work after restore.

- After editing, the entry appears as follows:

```
cat /root/.ssh/authorized_keys
no-port-forwarding,no-agent-forwarding,no-X11-forwarding,
ssh-rsa AAAAB3Nza...<truncated>..HruCzDsb3j
```

Target: Azure

- 1 AWS instances have Hv and NVMe drivers preinstalled and hence no additional steps are required. Confirm if the drivers exist in your instance, run the command:

```
lsinitrd | grep -i -e hv -e nvme
```

- 2 Replace the device names with **UUID** in `/etc/fstab` file.

SLES

Target: AWS

By default, AWS uses SUSE kernel. Hence install the SUSE kernel and select it at the boot time from the GRUB menu on the restored VM as follows:

- 1 Refer to the following documentation for more information on entries in `zypp.conf` file as it may affect the number of kernels retained and their behavior:

[Installing multiple kernel versions](#)

Proceed with the next steps after ensuring that the SUSE computer can work with multiple kernels.

- 2 Run the following command to list the available kernels:

```
zypper se -s 'kernel*'
```

- 3 From the list of kernels displayed in the above step, install a suitable default kernel version:

```
zypper in kernel-default-<VERSION>
```

For example, `zypper in kernel-default-5.3.18-53.3`

- 4 List the installed kernels and kernel modules using the following command:

```
zypper se -si 'kernel*'
```

- 5 It is recommended to set a root user password.

6 If the required drivers are not installed, then install them by running the following commands:

- `lsinitrd --kver <YOUR NEW KERNEL VERSION> | grep xen`
- Update/Create the `/etc/dracut.conf` file with the following line:
`add_drivers+="xen-blkfront xen-netfront nvme-core nvme"`
- Run the following command:
`dracut -f -v`
- `dracut -f -v --kver <YOUR NEW KERNEL VERSION>`
- `lsinitrd --kver <YOUR NEW KERNEL VERSION> | grep xen`

`<YOUR NEW KERNEL VERSION>` is the new kernel version installed in Step 3 above.

7 Back up the `/etc/default/grub` file. Edit the original `grub` file, add `GRUB_TIMEOUT` and `GRUB_TIMEOUT_STYLE` entries and comment out the following parameters:

```
GRUB_HIDDEN_TIMEOUT
GRUB_HIDDEN_TIMEOUT_QUIET
```

By default, `GRUB_DEFAULT` is set to 0 in the `/etc/default/grub` file. Change the default value so that it loads Azure kernel on restart and not the newly installed kernel.

- For example, `GRUB_DEFAULT='1>KERNEL_INDEX'` where `KERNEL_INDEX` can be found with `grub2-mkconfig` command or by analyzing `/boot/grub2/grub.cfg` file.
- Updating `GRUB_DEFAULT` ensures that the source VM keeps using Azure kernel, in an event, where it is restarted while the new kernel is installed.
- The GRUB config file has entries similar to the following:

```
#GRUB_HIDDEN_TIMEOUT=
#GRUB_HIDDEN_TIMEOUT_QUIET=true
GRUB_DEFAULT'1><YOUR KERNEL INDEX NUMBER>'
GRUB_TIMEOUT=20
GRUB_TIMEOUT_STYLE=menu
```

- Update the GRUB config file using the following command:
`grub2-mkconfig -o /boot/grub2/grub.cfg`
- After restore, to access GRUB menu during restart, press `ESC` twice during the countdown on EC2 Serial Console.

For more information on GRUB entries, refer to [Simple configuration handling](#)

Target: Azure

- 1 Run the following command to verify if the drivers are preinstalled:

```
lsinitrd | grep -i -e hv -e nvme
```

- 2 It is recommended to replace the device names with **UUID** in `/etc/fstab` file.

Ubuntu

Target: AWS

- 1 Run the following command to install the **linux-aws** kernel package:

```
sudo apt-get update && sudo apt-get upgrade -y linux-aws
```

- 2 Change grub countdown style and increase time-out in `/etc/default/grub` file. This enables the user to enter the recovery mode if issues are faced during the restart:

```
GRUB_TIMEOUT_STYLE=menu
```

```
GRUB_TIMEOUT=20
```

- 3 To avoid loading the new kernel on restart, ensure that the default kernel entry (GRUB_DEFAULT) in the grub configuration file is pointing to Azure-specific kernel and not the newly installed kernel.

- 4 Run the following command to update `grub` file:

```
update-grub
```

Target: Azure

- 1 Run the following command to verify if the drivers are preinstalled:

```
lsinitrd | grep -i -e hv -e nvme
```

- 2 It is recommended to replace the device names with **UUID** in `/etc/fstab` file.

Windows

Target: AWS

- 1 Check the boot mode (Legacy or UEFI).

- 2 Run the following:

```
(Ctrl + R) -> MSInfo32.exe -> BIOS Mode
```

- 3 Install the following drivers:
 - PV driver: [AWSPVDriver.zip](#)
 - EC2 install: [EC2Install.zip](#)

- NVME: [AWSNVMe.zip](#)
- ENA: [AwsEnaNetworkDriver.zip](#)

Post-restore configurations

This section describes the additional configurations required for the target VM on different operating systems to ensure that the restored instance is reachable.

RHEL

Target: AWS

- 1 Perform the following steps to resolve the issue of restored AWS instance check fail and the instance is not reachable:
 - Connect to EC2 Serial Console, check if the network interface is up, and VM has started with all mount points.
If you get the following failure message while accessing EC2 Serial Console, change the instance type of VM to the one which supports EC2 Serial Console:

```
This instance type is not supported for the EC2 serial console.
```
 - You might be logged into rescue mode if some critical mount points fail:
 - Check if the following articles are applicable:
 - [XFS reports "Log inconsistent \(didn't find previous header\)" when trying to mount filesystem](#)
 - [Unable to mount or check XFS filesystem](#)
 - Check if all partitions are mountable, if not, further troubleshooting might be required to find out why they are failing.
 - Drives with mount failures can be detached from the impaired VM and attached to a healthy instance in the same availability zone for further troubleshooting.
 - If mounting succeeds and there are no failures, try restarting before the next steps, till you can boot into user space.
- 2 If the VM is configured with `PasswordAuthentication` password check entry in `/etc/ssh/sshd_config` file, then set the value to **Yes**.
- 3 If it is configured with `PubkeyAuthentication` ssh key, then set the value of the entry in `/etc/ssh/sshd_config` file to **Yes**.

- 4 If the instance is unreachable and the following error messages are displayed when attempting to connect to the EC2 console:

```
[ERROR]: Azure datasource failure occurred: result=error |
reason=failure to find primary DHCP interface | agent=Cloud-Init
```

```
[ERROR]: Azure datasource failure occurred: result=error |
reason=http error 401 querying IMDS
```

It is recommended to create an empty file named `cloud-init.disabled` in the `/etc/cloud/` directory by executing the following command:

```
# touch /etc/cloud/cloud-init.disabled
```

This issue arises because the Cloud-init auto-configuration process during the VM boot sequence enters an infinite loop when it is unable to reach the required endpoints, resulting in a failure to successfully configure the endpoints.

- 5 Update the device names in `/etc/fstab` file if they were not replaced with UUID during pre-backup.

For example, change `/dev/sdc1` → `/dev/nvme1n1p1`. This allows to establish the network connection.

- 6 Now you can troubleshoot any mount failures and logon by ssh from any non-CSP Serial Console.

Target: Azure

To backup the restored Azure VM, deploy the Azure VM extension before taking the backup.

SLES

Target: AWS

Perform the following steps to resolve the issue of restored AWS instance check fail and the instance is not reachable:

- 1 Connect to EC2 Serial Console, check if the network interface is up, and VM has started with all mount points.

If you get the following failure message while accessing EC2 Serial Console, change the instance type of VM to the one which supports EC2 Serial Console:

```
This instance type is not supported for the EC2 serial console.
```

- 2 If the VM has booted into emergency mode further troubleshooting might be required.

- 3 Try mounting all partitions with `mount -a` command. If it fails, check if the following article is applicable:
[How to perform filesystem check \(fsck\) in the rescue mode?](#)
- 4 While restarting, the `cloud-init` service might cause the VM to get stuck in network configuration. Reconfigure or stop, disable, and uninstall/remove `cloud-init` service using the following respective commands:

```
systemctl stop cloud-init  
systemctl disable cloud-init  
zypper remove cloud-init
```

If required, reinstall `cloud-init` service after a restart.
- 5 Restart the VM.
- 6 During restart select AWS kernel from the GRUB menu.
- 7 Check the mount points.
- 8 Check the status of the network interface with `ip addr` command. If it is up you should be able to log on by SSH externally, else try SSH troubleshooting.

Target: Azure

To backup the restored Azure VM, deploy the Azure VM extension before taking the backup.

Ubuntu

Target: AWS

Perform the following steps to resolve the issue of restored AWS instance check fail and the instance is not reachable.

- 1 Stop and Start the VM.
- 2 Log on using EC2 Serial Console and switch to superuser.
- 3 Disable Azure-specific services if any:

```
systemctl disable hv-kvp-daemon.service  
systemctl disable walinuxagent.service  
systemctl disable walinux-agent
```
- 4 Reconfigure the data source for `cloud-init` service to point it to EC2 data source, and disable any Azure cloud-specific configurations:
 - `dpkg-reconfigure cloud-init`
 - Few configuration files can be moved or renamed.

For example,

```
mv /etc/cloud/cloud.cfg.d/10-azure-kvp.cfg  
/etc/cloud/cloud.cfg.d/10-azure-kvp.cfg.disabled  
mv /etc/cloud/cloud.cfg.d/90-azure.cfg  
/etc/cloud/cloud.cfg.d/90-azure.cfg.disabled
```

- `cloud-init clean --logs`

- 5 Set the default kernel to AWS by changing the value of **GRUB_DEFAULT** entry in `/etc/default/grub` file:

By default we have **GRUB_DEFAULT=0** in `/etc/default/grub` file. Change it to: **GRUB_DEFAULT='KERNEL_INDEX'** where the value of **KERNEL_INDEX** can be found with `update-grub` command.

- 6 Restart the VM.
- 7 If there is any issue after restart, enter the recovery mode for AWS kernel and drop to root shell prompt to perform any troubleshooting steps.

Target: Azure

To backup the restored Azure VM, deploy the Azure VM extension before taking the backup.

Windows

Target: AWS

Log on with the same username and password as the source VM.

Target: Azure

To backup the restored Azure VM, deploy the Azure VM extension before taking the backup.

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 88. 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 30.

See “[Protecting AWS or Azure VMs for recovering to VMware](#)” on page 88. 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 173.
- 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 173.
- 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 173.
- 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 173.
- 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 173.
- 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 173.
- 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 143.
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 Redshift clusters

You can restore Redshift clusters from the **PaaS** tab, under Cloud workload.

To restore Redshift clusters 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 In the **Recovery points** tab, click 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.
 - To restore to the image to the original location, click **Original location**, and then click **Start recovery**.
 - To restore to the image to an alternate location, click **Alternate location**. Select the required location from the list of available locations, and then click **Start recovery**.

Additional steps required after restoring a Redshift cluster

These additional steps are required because, even though the restore is successful, NetBackup may not restore one or more properties or attributes of the instance.

You can perform the following steps after restoring a Redshift cluster instance.

- (Optional) The attribute `publicallyaccessible` is set to `False`. You can manually set it to `True` from your AWS console.
- (Optional) The attribute `ClusterParameterGroupName` is not restored. You can manually configure it from your AWS console.

Recovering AWS DocumentDB and Neptune assets

You can restore AWS DocumentDB and Neptune assets from the **PaaS** tab, under Cloud workload.

To restore AWS DocumentDB and Neptune assets 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 In the **Recovery points** tab, click 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.
 - To restore to the image to the original location, click **Original location**, and then click **Start recovery**.
 - To restore to the image to an alternate location, click **Alternate location**. Select the required location from the list of available locations, and then click **Start recovery**.

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://support.cohesity.com/s/article/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 143.

(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.

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.

To restore a database to an elastic pool within the selected target instance, select **Use Elastic pool** and enter the name of an existing elastic pool.

- 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 node in the web UI, manually duplicate the image to an MSDP storage. See *NetBackup Web UI 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 178.

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.

Note: If a BFS (Backup from snapshot) + GRT (Granular restore) protection plan is run on an unconnected VM which has provider managed consistency enabled, then Single File Restore (SFR) is available only from the backup copy.

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

Table 4-1 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. |
| | Perform granular restoration of files and folders. |

Supported environment list

The following table lists the supported versions.

Table 4-2 Supported versions

| Application | Version |
|---|---|
| NetBackup | 11.2 |
| NetBackup backup host OS | <ul style="list-style-type: none"> ■ RHEL 8.8 onwards ■ Windows 16, 19 and 22 ■ OEL 8.8 onwards ■ SUSE Linux |
| Snapshot Manager host OS | <ul style="list-style-type: none"> ■ RHEL 8.6 onwards ■ SLES 15 ■ OEL 8.x and later ■ Ubuntu 18.04 LTS, 20.04 LTS, 22.04 LTS, and 24.04 LTS <p>Note: The version of the OS (Ubuntu 20.04 LTS) listed on the UI is the version of the container.</p> |
| Cloud providers | <ul style="list-style-type: none"> ■ Amazon Web Services ■ Microsoft Azure ■ Microsoft Azure Stack Hub ■ Google Cloud Platform ■ Oracle Cloud Infrastructure. Oracle PCA is not supported. |
| 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 ■ Oracle Cloud Infrastructure: VM.Standard.E4.Flex/ VM.Standard.E5.Flex/ VM.Standard3.Flex/ VM.Optimized3.Flex |
| Snapshot Manager agent host to be protected | <ul style="list-style-type: none"> ■ Linux OS: RHEL 8.8 onwards, OEL 8.x or later ■ 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) 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"> ■ ext3 ■ ext4 ■ xfs | <ul style="list-style-type: none"> ■ GPT ■ MBR ■ No layout (direct FS) |
| Windows (With consistent snapshot property) | NTFS | <ul style="list-style-type: none"> ■ GPT ■ MBR |

Note: 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.
 - (OCI, and GCP):The Snapshot Manager host and the connected VM must be in the same tenancy/project and region.

- (OCI): The Block volume management plug-in must be enabled on the Snapshot Manager host, as well as the connected VMs.
- (OCI): Ensure that Oracle cloud agent is installed and active on Snapshot Manager and the protected VMs.
- 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 189.
- Configure ports 5671 and 443 for open Snapshot Manager host.
- For agentless restore, in Linux and Windows systems, configure port 22 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 `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.

- Until the old agent (preinstalled) 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.
- Before running the protection plan created for granular restore on Virtual machines on cloud, ensure that you regenerate and set the disks' UUIDs when disks created from same snapshot are attached to the Virtual Machine on cloud. The `/etc/Estab` entry must also be updated with the new UUID.
- Granular restore 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 Hub, AWS, OCI, and GCP. However, VxMS indexing is not supported for volumes or partitions created on software RAID devices. These volumes, or partitions, are skipped while indexing the file system.
- 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.

- While restoring application or file systems from one OS version to another OS version, refer to the OS and application vendor's compatibility matrix. The restore of a 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.
- Consider a scenario where an alternate path that does not exist on the RHEL target host is specified for a 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 VMs having VHDX disks (Azure Ultra disks, Premium SSD v2 with a 4k sector size) and granular restore.
- When snapshots are performed or indexed, the following devices are ignored:
 - Ephemeral storage devices: For example, an Amazon AWS instance store volumes and Microsoft Azure temporary disks

Note: These devices are also ignored for indexing also.

- File systems that are created on the LDM disk.

Note: Though the files/folders from the LDM disks are visible for selecting on Web UI during Single File Restore, the files would not be restored and the restore job would fail.

- Extended attributes would not be restored for Linux VM files.
- For FIPS setup, Single file restore for Windows to Windows is not supported.

- Linux VM single file restore: If a directory contains more than 100K files, restore of the directory and files inside the directory would be skipped due to a limitation in Instant Access mount.

Limitations for single file restore in OCI

- You must attach the block volumes after creating the VMs, and provide the consistent device path while attaching the volume.
- Granular restore from a snapshot copy and volume restore, require the block volume management plug-in to be enabled on the target VM. Restart the VM after enabling the plug-in.
- For Windows instances, granular restore from a snapshot copy is not supported.
- Granular restore from backup copy to an Windows instance, you need to manually copy the granules from the NFS share.
- Granular restore from a snapshot copy, from a boot volume to another boot volume is not supported for all operating systems.
- Granular restore from Linux OS not having consistent device paths are not supported, if the disks are attached as paravirtualized attachments.
- Granular restore from a source VM that has higher kernel version, to a target VM with lower kernel version is not supported.

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 Cohesity 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 or later, then the primary server on version 10.3 or later 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, GCP, OCI, and AWS, NetBackup supports snapshot and recovery of cloud assets that are encrypted using the keys that the manager provides.

Files and folders are restored more quickly on target virtual machines (VMs) utilizing OnHost agents compared to those using an agentless approach. Moreover the restoration performance is inversely proportional to the number of disks already attached to the target VM.

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 | <p>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.</p> | <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 Hub 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. |

Dialog box **Snapshot copy** **Backup copy**

- Restore target options**
- **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**
- **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.

| Dialog box | Snapshot copy | Backup copy |
|-------------------------|---------------|---|
| 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. <p>Note: Only private IPs or host names are allowed.</p> |

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.

- To retain the hardlink between source and its linked files, select source file and linked files during restore and uncheck the **Create new files for hard links** checkbox.

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 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: For OCI, the Block volume management plug-in must be enabled on the client VMs to restore volumes.

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"> ■ For agentless: <code>/opt/VRTScloudpoint/.agent/flexsnap-agentless-onhost.log</code> ■ For on-host agent: <code>/var/log/flexsnap/flexsnap-agentless-onhost.log</code> ■ For single file restore specific logs, use the path provided in the Activity monitor. | <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) | 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: <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 | The restore job fails with the following error when restoring files and folders to the target VM which is deleted from the cloud environment: <pre>Error nbcs (pid=44859) Target VM not found, asset_id <asset_id></pre> | Select a different target VM. |
| Creating instant access fails | If instant access is not enabled on the MSDP storage server, the creation of the instant access fails during the restore job. | Verify if instant access is supported on the MSDP media server. Run the following pre-check script: <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 | 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. | Select a smaller number of volumes for the restore. |
| Not enough space: **\driverMapping.json | The media server where MSDP is configured has FIPS enabled. | Disable FIPS on the media server where MSDP is installed. Or, add the domain user Samba credentials to the target VM. |

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.

Issues in snapshot restore from OCI

Invalid target path. Both the source and destination disks are boot disks, or the disks are mounted on

```
/dev/oracleoci/oraclevda
```

This error occurs when you try granular restore from a snapshot with at least one of the selected files from the boot volume, and the destination volume of restoration being a boot volume.

Workaround:

To restore files and folders from the boot volume or file system, specify a block volume as the target path.

The selected files/folders were skipped from recovery as the snapshot does not contain the corresponding file system.

During granular restore from snapshot, this error occurs for VMs that do not have the Oracle Cloud Agent installed, or have the disks attached using the paravirtualized attachment method.

Error in the Activity monitor:

```
[{'error': 'The selected files/folders were skipped from recovery as  
the snapshot does not contain the corresponding file system.',  
'mount': 'Unknown Mount Point/Drive'}}]
```

Workaround:

Do the following:

- Install the Oracle Cloud Agent if it is supported on the platform by Oracle. Or, attach all the disks using iSCSI attachment type, including the boot disk.
- Wait for the discovery to complete and retry the backup.

Granular restore from snapshot copy fails with the error “Block Volume Plug-in: must be enabled on the instance...”

This error occurs when the Block Volume Management plug-in for the Oracle Cloud Agent is not enabled on the target VM.

Workaround:

Do the following:

1. Enable the Block Volume Management plug-in from the OCI Console, in the target VM.
2. Restart the VM.
3. Try to restore again.

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>](#)
- [VMs and other OCI assets with CMK-encrypted disks are marked as deleted in NetBackup UI.](#)
- [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](#)
- [Status Code 150: Termination requested by administrator](#)
- [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 12.

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 5-1 Log files for configuration

| Process | Logs |
|--|--|
| <p>tpconfig</p> <p>tpconfig command is one way to register Snapshot Manager in NetBackup.</p> | <p>Windows</p> <p><i>NetBackup install path</i>\Volmgr\bin\tpconfig.exe</p> <p>UNIX</p> <p>/usr/opensv/volmgr/bin/tpconfig</p> |
| <p>nbwebservice</p> <p>Plug-ins are configured using NetBackup REST API.</p> | <p>Windows</p> <p><i>NetBackup install path</i>\NetBackup\wmc\webserver\logs</p> <p>UNIX</p> <p>/usr/opensv/wmc/webserver/logs</p> <p>/usr/opensv/logs/nbwebservices</p> |
| <p>nbemm</p> <p>nbemm stores the Snapshot Manager and plug-in information in the EMM database.</p> | <p>Windows</p> <p><i>NetBackup install path</i>\NetBackup\logs\nbemmm</p> <p>UNIX</p> <p>/usr/opensv/logs/nbemmm</p> |

Log files for asset discovery

Use the following logs to troubleshoot asset discovery issues.

Table 5-2 Log files for asset discovery

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

Log files for snapshot creation

Use the following logs to troubleshoot snapshot creation issues.

Table 5-3 Log files for snapshot creation

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

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

| Process | Logs |
|---|--|
| <p>nbc</p> <p>nbc PID for a given job is available in the NetBackup activity monitor.</p> | <p>Windows</p> <pre>NetBackup install path/bin/vxlogview -i 366 -P nbc_process_id</pre> <p>UNIX</p> <pre>/usr/opensv/netbackup/bin/vxlogview -i 366 -P nbc_process_id</pre> <p>The nbc logs are available at the following location:</p> <p>Windows</p> <pre>NetBackup install path/logs/ncfnbc</pre> <p>UNIX</p> <pre>/usr/opensv/logs/ncfnbc</pre> |
| <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> <pre>NetBackup install path/bin/vxlogview -o 118</pre> <p>UNIX</p> <pre>/usr/opensv/netbackup/bin/vxlogview -i 118</pre> |

Log files for restore operations

Use the following logs to troubleshoot restore issues.

Table 5-4

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

Table 5-4 (continued)

| Process | Logs |
|---|--|
| ncfnbcs nbcs PID for a given job is available in the NetBackup activity monitor. | Windows <i>NetBackup install path/bin/vxlogview -i 366 -P nbcs_process_id</i> UNIX <i>/usr/opensv/netbackup/bin/vxlogview -i 366 -P nbcs_process_id</i> |

Log files for snapshot deletion

Use the following logs to troubleshoot snapshot deletion issues.

Table 5-5 Log files for snapshot deletion

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

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

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

This may happen when you have 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.

VMs and other OCI assets with CMK-encrypted disks are marked as deleted in NetBackup UI.

Explanation:

If the KMS service at the OCI provider is down, the VMs and other assets with CMK-encrypted disks are marked as deleted in NetBackup UI. Once the KMS service is restored, the deleted status is cleared after a successful plug-in level discovery, and the assets or VMs become eligible for backup. No further action is required.

Workaround:

Ensure that the KMS service at the OCI provider-end is running.

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.

Status Code 150: Termination requested by administrator

Explanation: This appears when you manually cancel a backup, a snapshot, or a restore job from the activity monitor.

In case of a restore job, a virtual machine or volume is created on the portal during the restore operation. A restore job may not have the associated NetBackup job because it is canceled. If the resources are created in the cloud from the cloud cost perspective, the cloud administrator must review the newly created resource.

Workaround: Manually clean up the virtual machine or volume from the provider portal. Also clean up any temporary staging area volumes that are created.

Troubleshoot PaaS workload protection and recovery issues

Backup job fails or falls back to full for a PaaS database

If you manually delete a backup image created by an incremental schedule, the next incremental backup may either fail or run as a full backup. The outcome depends on which image was removed.

For example, consider a backup sequence with images F1 (full), I1, I2, and I3 (incremental).

- If the most recent incremental image (I3) is deleted, the next backup automatically switches to a full backup.
- If an earlier incremental image (such as I1 or I2) is deleted, future incremental backups fail. The error indicates missing images between the last full backup and the current backup time.

The error message specifies the expected number of images (X) and the actual number found (Y) in the catalog.

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>' '<objctId>' does not have authorization to perform action 'Microsoft.Sql/servers/databases/read' over scope '<resoourceld>' or the scope is invalid. If access was recently granted, 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 `bpdbm` 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.

EMM status: Unsupported media server.

Explanation: This occurs when the media server is running a lower version of NetBackup or is hosted on an unsupported platform.

Workaround: You must have at least one media server with the following configuration:

- NetBackup version 11.1 or later.
- Hosted on Red Hat Enterprise Linux

You can either add a new media server with the supported version and platform, or upgrade an existing one to meet these requirements.

An unexpected failure occurred

Explanation: Job details include the following error:

```
An unexpected failure occurred: Data plan execution failed with the
message - "One or more errors occurred. (One or more errors occurred.
(One or more errors occurred. (Exception of type
'System.OutOfMemoryException' was thrown.)")
```

Workaround: This issue is caused by a limitation in Microsoft's `SqlPackage` utility.

To resolve it, increase the available memory on the media server to match the database size, and then retry the job.

There is no exact formula to calculate the required memory, but it is recommended to allocate at least twice the memory of the database schema size.

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 the **Activity monitor**.

The bpdbm log confirms the same by displaying the following error message:

```
<16> db_error_add_to_file: cannot stat(\\?\C:\Program Files\Cohesity
NetBackup\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_FULLL.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_FULLL.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_FULLL.f_imgUserGroupNames0
```

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

Workaround: Use a policy prefix name in the 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: When using system or user-managed identity for the PaaS Azure SQL and Managed Instances, apply the same permissions and 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.

Backup job fails with the error: Invalid object name 'cdc.<schema_name>_<table_name>_CT'

Explanation: Backup failed due to an error while executing the backup query script on the Azure SQL Server. The following error message appears:

```
('42S02', "[42S02] [Microsoft][ODBC Driver 18 for SQL Server] [SQL Server]Invalid object name 'cdc.<schema_name>_<table_name>_CT'. (208) (SQLExecDirectW)")
```

This issue may occur due to a mismatch between user tables and their associated CDC tables. The mismatch typically arises after schema changes or when database objects are moved between schemas after CDC is enabled.

Recommended action:

Do the following:

1 Disable CDC on the database before running the backup. Use the command:

```
# EXEC sys.sp_cdc_disable_db
```

2 Run a full backup after disabling CDC.

NetBackup automatically re-enables CDC during the next backup after you manually disable it.

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.

Importing from a replica on a target NetBackup domain fails, with the status code 191.

Explanation: The import operation on the target domain may fail with the state code 191: `No images successfully processed`. The job details in the Activity monitor show: `Failed to create the JSON payload`.

Cause: The image that you are replicating to the target domain is created from a NetBackup 10.4 or older media server, which does not have the required metadata in the NetBackup catalog.

Workaround: Do any of the following:

- Use a media server version later than 10.4 to use the AIR feature for the PaaS workloads.
- Install EEB on a 10.4 media server to use this as a back-level media server for the AIR feature for PaaS workloads. Contact Cohesity Technical Support for more details.

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.

Troubleshooting Azure Postgres issues

Incremental backup job fails with the error: Cannot create the replication slot as the maximum limit is reached

Explanation:

The number of replication slots created on the server has exceeded the configured `max_replication_slot` server parameter.

Workaround:

Do any of the following:

- Delete the unused replication slots.
- Increase the value of `max_replication_slots` in the server parameters.

Backup fails with an error: Could not write to data file for XID 1676198: No space left on the device

Explanation:

WAL is full and it reached the configured value in `WAL_SIZE` server parameter.

Workaround:

Increase the value of `WAL_SIZE` in server parameters.

Troubleshooting Amazon RDS Custom for SQL issues

Access denied error for RDS Custom SQL on-premises agent.

Explanation:

This error occurs in post-backup operations while deleting the batch files created during the backup operation. You can manually delete these batch files when there is no backup job running for that instance.