NetBackup™ Web UI Cloud Object Store Administrator's Guide

Release 10.1.1

VERITAS™

NetBackup™ Web UI Cloud Object Store Administrato's Guide

Last updated: 2022-12-19

Legal Notice

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

http://www.veritas.com

Technical Support

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

https://www.veritas.com/support

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

https://my.veritas.com

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

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

Documentation

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

https://sort.veritas.com/documents

Documentation feedback

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

NB.docs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1

Introduction

This chapter includes the following topics:

Overview of NetBackup protection for Cloud object store

The NetBackup web UI provides the capability for backup and restore of Cloud object stores. You can deploy the NetBackup environment in the same cloud network as the object store. Alternatively, you can provide http(s) connectivity to the object store service endpoint and the backup host. You can deploy NetBackup outside the cloud vendor as well.

Note: Cloud vendors may levy substantial charges for data egress for moving data out of their network. Check your cloud provider pricing for data-out before configuring a backup policy that transfers data out of one cloud to another cloud region or an on-premises data center.

NetBackup can protect Azure Blob Storage, and a wide variety of S3 API-compatible object stores like AWS S3, Google Cloud Storage (GCS), Hitachi Cloud Platform object store, and so on. For a complete list of compatible object stores, refer to the NetBackup Hardware Compatibility List (HCL).

Features of NetBackup Cloud object store workload support

Table 1-1 Salient features

Feature	Description
Integration with NetBackup role-based access control (RBAC)	The NetBackup web UI provides the **Default cloud object store Administrator** RBAC role to control which NetBackup users can manage Cloud object store operations in NetBackup. The user does not need to be a NetBackup administrator to manage Cloud object store operations.
Management of Cloud object store accounts	You can configure a single NetBackup primary server for multiple Cloud object store accounts, across different cloud vendors as required.
Authentication and credentials	Wide emphasis for security. For protecting Azure Blob Storage, Storage account, and Access Key must be specified. For all S3 API-compliant cloud vendors, Access key and Secret Key are supported. For Amazon S3, in addition to Access Key, IAM role and Assume role mechanism of authentication are also supported.
Backup policy	A single backup policy can protect multiple S3 buckets or Azure blob containers from one Cloud object store account.
Intelligent selection of cloud objects	Within a single policy, NetBackup provides flexibility to configure different queries for different buckets or containers. Some buckets or containers can be configured to backup all objects in them. You can also configure some bucket and containers with intelligent queries to identify objects based on: ■ Object name prefix ■ Entire object name ■ Object tags

Table 1-1 Salient features *(continued)*

Feature	Description
Scalable solution	In addition to full backup, NetBackup also supports different types of incremental schedules for faster backups. Accelerator feature is also supported for Cloud object store policy. Enable checkpoint restart in the policy to be able to restart a failed or suspended job, right from the point that it stopped. We do not need to repeat the entire data transfer from the start of the job.
Granular restore	NetBackup supports an easy way to restore all objects in a bucket/container, as well as provides prefix, folder and object-based views to restore only a selected subset of the objects. You can narrow down a selection of backup images for restore in NetBackup by providing a date and time range.
Restore options	NetBackup supports adding an arbitrary prefix to all objects when restoring. Thereby, restores the objects with a different name when you do not want the restored objects to interfere with the original objects. By default, NetBackup skips overwriting objects that already exist in the cloud object store to conserve on bandwidth and cloud costs. You can change this default behavior using the Overwrite option, so that restored copies can overwrite the cloud object store copies.
Alternate location restore	Objects selected for restore can be restored: ■ To the same bucket or container ■ To a different bucket or container in same account ■ To an altogether different cloud account of the same cloud vendor.

Chapter 2

Managing Cloud object store assets

This chapter includes the following topics:

Prerequisites for adding Cloud object store accounts

Gather the following before you start adding a Cloud object store account.

- Gather information about the cloud provider, service host, and region.
- Check the supported authentication types by the cloud service provider and decide on the authentication type to use. All cloud providers support the Access credentials authentication type. Other supported Authentication types are:
  - IAM Role (EC2): For Amazon and Amazon Gov
  - Assume Role: For Amazon and Amazon Gov
  - Credential Broker: For Amazon Gov
- If you plan to use proxy for communication with cloud endpoints, gather the required details of the endpoints.
- Get the Cloud account credentials, and any additional required parameter, as per the authentication type. These credential details should have required

permissions recommended in NetBackup documentation. See "Permissions required for Amazon S3 cloud provider user" on page 10.

- Make sure that the required ports are open, and configurations are done for communication from the backup host to cloud provider endpoint using REST API calls.

Permissions required for Amazon S3 cloud provider user

Amazon (S3) cloud provider, requires the following permissions to work with NetBackup:

- s3:ListAllMyBuckets
- s3:ListBucket
- s3:GetBucketLocation
- s3:GetObject
- s3:PutObject
- s3:GetObjectTagging
- s3:GetObjectAcl
- s3:PutObjectAcl
- s3:PutObjectTagging
- s3:RestoreObject

Adding Cloud object store accounts

Adding a Cloud object store account is the first step in protecting a workload. You can add as many accounts as required. You can create different Cloud object store accounts to fit your business logic. For example, grouping buckets from a particular cloud service provider. AWS S3 compatible accounts, require separate RBAC access rights for backup and restore. You can create separate accounts for backup and restore to better organize the access rights.

Depending on the bucket or container which you want to protect, you must add at least one Cloud object store account, per cloud service provider, per region.

You may need to create multiple Cloud object store accounts, for the same cloud service provider and region, to better organize settings like SSL, proxy, and the type of credentials to be used for the set of buckets or containers.

The required permissions for backup and recovery are different. See if it is helpful to create separate accounts for backup and recovery. You need to use other than original bucket option, to restore to a different Cloud object store account during recovery.

Note: Cloud object store account shares the namespace with Cloud storage server and MSDP-C LSU name.

For Cloud object store account, NetBackup supports a variety of cloud providers using AWS S3 compatible APIs (for example Amazon, Google, Hitachi etc.), other than Microsoft Azure. For such providers, you need to provide AWS S3 compatible account access details to add the credentials (that is, Access Key ID, Secret Access key) of the provider.

To add a Cloud object store account:

1. On the left, click **Cloud object store** under **Workloads**.
2. In the **Cloud object store account** tab, click **Add**.
3. Enter a name for the account in **Cloud object store name** field, select a provider from the list **Select Cloud object store provider**, and select a backup host from **Backup host for validation** list. Credential validation, backup, and recovery of the Cloud object stores are supported by NetBackup 10.1 or later on RHEL media server.
4. Select a region from the available list of regions. Click **Add** above the **Region** table to add a new region.

   See "Adding a new region" on page 18.. Region is not available for some Cloud object store providers.

   For GCP, which supports dual region buckets, select the base region during account creation. For example, if a dual region bucket is in the regions *US-CENTRAL1*, *US-WEST1*, select *US*, as region during account creation to list the bucket.
5. In **Access settings** page: Select a type of access method for the account:
   - **Access credentials**-In this method, NetBackup uses the Access key ID, and the secret access key to access and secure the Cloud object store account. If you select this method, perform the subsequent steps 6 to 10 as required to create the account.
   - **IAM role (EC2)**-NetBackup retrieves the IAM role name and the credentials that are associated with the EC2 instance. The selected backup host must be hosted on the EC2 instance. Make sure the IAM role associated with EC2 instance has required permissions to access the required cloud

resources for Cloud object store protection. Make sure that you select correct region as per permissions given to EC2 instance while configuring the Cloud object store account with this option. If you select this option, perform the optional steps 7 and 8 as required, and then perform steps 9 and 10.

- **Assume role**-In this method, NetBackup uses the provided key, the secret access key, and the role ARN to retrieve temporary credentials for the same account and cross account. Perform the steps 6 to10 as required to create the account.
  See "Creating cross account access in AWS " on page 14.
- **Credentials broker**- NetBackup retrieves the credentials to access the cloud resources required for Cloud object store protection.

6 You can add existing credentials or create new credentials for the account:

- To select an exiting credential for the account, select the **Select existing credentials** option, select the required credential from the table, and click **Next**.
- To add a new credential for the account, select **Add new credentials**. Enter a **Credential name**, **Tag**, and **Description** for the new credential.
  For cloud providers supported through AWS S3 compatible APIs, use **AWS S3 compatible** credentials. Specify the **Access key ID** and **Secret access key**.
  For Microsoft Azure cloud provider, provide **Azure Blob** credentials, specify **Storage account** and **Access key**.
- If you use **Assume role** as the access method, specify the Amazon Resource Name (ARN) of the role to use for the account, in the **Role ARN** field.

7 (Optional) Select **Use SSL** if you want to use the SSL (Secure Sockets Layer) protocol for user authentication or data transfer between NetBackup and cloud storage provider.

- **Authentication only**: Select this option, if you want to use SSL only at the time of authenticating users while they access the cloud storage.
- **Authentication and data transfer**: Select this option, if you want to use SSL to authenticate users and transfer the data from NetBackup to the cloud storage, along with user authentication.
- **Check certificate revocation (IPv6 not supported for this option)**: For all the cloud providers, NetBackup provides a capability to verify the SSL certificates against the CRL (Certificate Revocation List). If SSL is enabled and the CRL option is enabled, each non-self-signed SSL certificate is

verified against the CRL. If the certificate is revoked, NetBackup does not connect to the cloud provider.

Note: NetBackup supports only Certificate Authority (CA)-signed certificates while it communicates with cloud storage in the SSL mode. Ensure that the cloud server (public or private) has CA-signed certificate. If it does not have the CA-signed certificate, data transfer between NetBackup and cloud provider fails in the SSL mode.

Note: The FIPS region of Amazon GovCloud cloud provider (that is s3-fips-us-gov-west-1.amazonaws.com) supports only secured mode of communication. Therefore, if you disable the Use SSL option while you configure Amazon GovCloud cloud storage with the FIPS region, the configuration fails.

---

8 (Optional) Select the **Use proxy server** option to use proxy server and provide proxy server settings. Once you select the **Use proxy server** option, you can specify the following details:

- **Proxy host**–Specify IP address or name of the proxy server.
- **Proxy Port**–Specify port number of the proxy server.
- **Proxy type**– You can select one of the following proxy types:
  - **HTTP**

    Note: You need to provide the proxy credentials for HTTP proxy type.

  - **SOCKS**
  - **SOCKS4**
  - **SOCKS5**
  - **SOCKS4A**

Select **Use proxy tunneling** for HTTP proxy type.

After you enable **Use proxy tunneling**, HTTP CONNECT requests are sent from the backup or recovery host to the HTTP proxy server. The TCP connection is directly forwarded to the cloud back-end storage. The data passes through the proxy server without reading the headers or data from the connection.

Select one of the following authentication types if you use HTTP proxy type.

- **None**– Authentication is not enabled. User name and password are not required.
- **Basic**–Username and password needed.
- **NTLM**–Username and password needed.

User name is the username of the proxy server.

Password can be empty. You can use maximum 256 characters.

9 Click **Next**.

10 In the **Review** page, review the entire configuration of the account, and click **Finish** to save the account.

NetBackup creates the Cloud object store accounts only after validation of the associated credentials with the connection information provided. If you face an error, update the settings as per the error details. Also, check if the provided connection information and credentials are correct. The backup host that you assign for validation, can connect to cloud provider endpoints using the provided information.

Creating cross account access in AWS

If you have multiple AWS accounts in your environment, and NetBackup deployed in one account, can protect data in other accounts as well. You need to configure cross account data access in AWS portal, before selecting **Assume role** as your access method. NetBackup only needs the access key, secret key, and role ARN.

Follow the guidelines in AWS documentation for creating cross account access. Briefly, you need to perform the following steps.

To configure AWS cross accounts:

1 Log on to the AWS provider portal.

2 Create a new IAM role in the target AWS account, which you want to protect.

3 Create a new policy for the IAM role and ensure that it has required permissions to access the bucket and objects in that target AWS account. See "Permissions required for Amazon S3 cloud provider user" on page 10.

4 Establish a trust relationship between the source and the target AWS accounts.

5 In the source AWS account, create a policy that allows the IAM role in the source AWS account, to assume the IAM role in the target AWS account.

6 Attach the policy to the source account user, whose access key and secret access key you use for the assume role.

Check certificate for revocation

For all the cloud providers, NetBackup provides a capability to verify the revocation status of SSL certificates using Online Certificate Status Protocol (OCSP). If SSL and the **Check certificate revocation** option, both are enabled, NetBackup verifies each SSL certificate. To verify, NetBackup makes an OCSP request to the CA to check revocation status of certificate presented during SSL handshake. NetBackup does not connect to the cloud provider, if the status is returned as revoked, or it failed to connect to the OCSP endpoint present in the SSL certificate.

To enable validation, update the USE_CRL property from the Cloud object store account dialog.

Requirements for enabling 'Check certificate revocation' option

- OCSP endpoints are HTTP thus, turn off any firewall rule that block HTTP (port 80) connection to external network. For example, http://ocsp.sca1b.amazontrust.com
- OCSP URL is dynamically retrieved from the certificate thus, disable any firewall rule that blocks unknown URLs.
- Typically, OCSP URLs endpoint support IPV4. For IPV6 environments disable the 'Check certificate revocation' option.
- Private Clouds typically have a self-signed certificate. Thus, for private clouds, Check certificate revocation is not required. Disable this check while configuring the account, otherwise, account creation fails.
- OSCP URL of CA should be present in certificate's 'Authority Information Access' extension.

Managing Certification Authorities (CA) for NetBackup Cloud

NetBackup cloud supports only X.509 certificates in .PEM (Privacy-enhanced Electronic Mail) format.

You can find the details of the Certification Authorities (CAs) in the `cacert.pem` bundle at following location:

- Windows:
  `<installation-path>\NetBackup\var\global\cloud`
- UNIX:
  `/usr/openv/var/global/cloud/`

Note: In a cluster deployment, NetBackup database path points to the shared disk, which is accessible from the active node.

You can add or remove a CA from the `cacert.pem` bundle.

After you complete the changes, when you upgrade to a new version of NetBackup, the `cacert.pem` bundle is overwritten by the new bundle. All the entries that you may have added or removed are lost. As a best practice, keep a local copy of the edited `cacert.pem` file. You can use the local copy to override the upgraded file and restore your changes.

Note: Ensure that you do not change the file permission and ownership of the `cacert.pem` file.

To add a CA

You must get a CA certificate from the required cloud provider and update it in the cacert.pem file. The certificate must be in .PEM format.

1. Open the `cacert.pem` file.
2. Append the self-signed CA certificate on a new line and at the beginning or the end of the `cacert.pem` file.

   Add the following information block:

   ```
   Certificate Authority Name

   ==========================

   -----BEGIN CERTIFICATE-----

   <Certificate content>

   -----END CERTIFICATE-----
   ```

3. Save the file.

To remove a CA

Before you remove a CA from the `cacert.pem` file, ensure that none of the cloud jobs are using the related certificate.

1. Open the `cacert.pem` file.
2. Remove the required CA. Remove the following information block:

```
Certificate Authority Name

==========================

-----BEGIN CERTIFICATE-----

<Certificate content>

-----END CERTIFICATE-----
```

3. Save the file.

List of CAs approved by NetBackup

- Baltimore CyberTrust Root
- Cybertrust Global Root
- DigiCert Assured ID Root CA
- DigiCert Assured ID Root G2
- DigiCert Assured ID Root G3
- DigiCert Global Root CA
- DigiCert Global Root G2
- DigiCert Global Root G3
- DigiCert High Assurance EV Root CA
- DigiCert Trusted Root G4
- GeoTrust Global CA
- GeoTrust Primary Certification Authority
- GeoTrust Primary Certification Authority - G2
- GeoTrust Primary Certification Authority - G3
- GeoTrust Universal CA
- GeoTrust Universal CA 2
- RSA Security 2048 v3
- Starfield Services Root Certificate Authority - G2
- Thawte Primary Root CA
- Thawte Primary Root CA - G2

- Thawte Primary Root CA - G3
- VeriSign Class 1 Public Primary Certification Authority - G3
- VeriSign Class 2 Public Primary Certification Authority - G3
- Verisign Class 3 Public Primary Certification Authority - G3
- VeriSign Class 3 Public Primary Certification Authority - G4
- VeriSign Class 3 Public Primary Certification Authority - G5
- VeriSign Universal Root Certification Authority

Adding a new region

You can add new regions for specific cloud object store account that you create in NetBackup, while creating the account. When you add a region, access is restricted to the specified region. Region is not available for some Cloud object store providers.

To add a region:

1. Enter a unique region name, and for **Location constraint**, enter the location identifier that the cloud provider service uses to access the buckets/containers of the associated region. For a public cloud storage, you need to get the location constraint details from the cloud provider.

   For the cloud providers that supports AWS v4 signature, specifying the **Location constraint** field is mandatory. You can obtain the correct value of the location constraint by using the `getBucketLocation` API on the concerned bucket. If this API returns location constraint as blank, use 'us-east-1' as the location constraint.
2. Enter the service URL. For example:
   `https://hostname:port_number/service_path`
3. Select the endpoint access style for the cloud service provider. If your cloud service provider additionally supports virtual hosting of URLs, select **Virtual Hosted Style**, otherwise select **Path Style**.
4. Specify the HTTP and HTTPS ports to use for the region.
5. Click **Add**. The added region appears in the **Region** table in the **Basic properties** page.

Manage Cloud object store accounts

The **Cloud object store** tab lets you view, add, edit, and delete the Cloud object store accounts. You can also validate the credentials for a Cloud object store account from this tab.

To view Cloud object store accounts

1. On the left, click **Workloads > Cloud object store**.
2. On the **Cloud object store account** tab, the accounts that are available to you are displayed.

Edit a Cloud object store account

You cannot update the provider, the selected service host, or the region on the **Edit** page.

To change the region, you may have to delete and recreate the Cloud object store account. You can do it in a maintenance window when the account is active and no job is associated with it. You can also update the region in **Host properties > Cloud storage** for the primary server.

To edit a Cloud object store account

1. On the left, click **Workloads > Cloud object store**.
2. Select the account that you want to edit. Then click **Edit**.

   See "Adding Cloud object store accounts" on page 10.

Validate credentials for a Cloud object store account

To validate credentials for a Cloud object store account

1. On the left, click **Workloads > Cloud object store**.
2. Select the account that you want to edit. Then click **Validate**.

   The result of the validation process is displayed in the same column.

Delete a cloud object store account

If you delete a Cloud object store account, NetBackup no longer protects any policies that are associated with this account. You can still recover existing backup images, using a different Cloud object store account. The backups of the policies that are associated with this Cloud object store account fail.

To delete a cloud object store account

1. On the left, click **Workloads > Cloud object store**.
2. Select the account that you want to edit. Then click **Delete**.
3. Click **Delete**.

Chapter 3

Protecting Cloud object store assets

This chapter includes the following topics:

About accelerator support

NetBackup accelerator for Cloud object store increases the speed of backups. The increase in speed is made possible by change detection techniques on the backup host. The backup host uses the change detection techniques to determine the current state of Cloud object store's objects or blobs to identify the changes that occurred since the last backup. The backup host sends the changed data to the media server in a more efficient backup stream. The media server combines the changed data with the rest of the cloud object store data that is stored in previous backups. If a portion of an object or blob is already in storage and has not been changed, the media server uses the copy in storage rather than reading it from the client. Accelerator backup has the following advantages:

- Reduces the I/O and CPU overhead on the client. The result is a faster backup and less load on the client.
- Creates a compact backup stream that uses less network bandwidth between the backup host and media server.
- Creates a backup image that contains all data that is needed for restore.

How NetBackup accelerator works with Cloud object store

The NetBackup accelerator creates the backup stream and backup image as follows:

- If the backup host has no track log for the given policy, bucket and query, NetBackup performs a full backup and creates a track log. The track log contains information about the objects/blobs data which is backed up as per query criteria, for comparison at the next backup.
- At the next backup, NetBackup identifies data and/or metadata that has changed since the previous backup. To do so, it compares information from the track log against information from the Cloud object store for each object/blob as per the query criteria for the bucket.
- The NetBackup backup host sends the following stream to the media server: The object/blobs' changed blocks, and the previous backup ID and data extents (block offset and size) of the unchanged blocks.
- The media server receives the object/blobs' changed blocks and the backup ID and data extents of the unchanged blocks. From the backup ID and object/blob descriptors, the media server locates the rest of the object/blob's data in existing backups.
- The media server directs the storage server to write the changed blocks, and combine these blocks with the locally stored, previously unchanged blocks to make a new full image."

Accelerator notes and requirements

Note the following about the NetBackup accelerator:

- NetBackup accelerator must be properly licensed. For the latest information on licensing, contact your NetBackup sales or partner representative.
- Supports the disk storage units only. Supported storage includes Media Server Deduplication Pool, NetBackup appliance, cloud storage, and qualified third-party OST storage. For supported storage types, see the NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List at the following URL: http://www.netbackup.com/compatibility
- Storage unit groups are supported only if the storage unit selection in the group is Failover.
- Supports the full backups and incremental backups.
- For every policy that enables the Use Accelerator option, the following backup schedules are recommended at a minimum: A full backup schedule with the Accelerator forced rescan option enabled. Another full backup schedule without the Accelerator forced rescan option enabled. See "Accelerator force rescan for Cloud object store (schedule attribute)" on page 24.
- If a previous backup of the policy, bucket and query does not exist on the backup host, NetBackup performs a full backup, and creates a track log on the backup host. This initial backup occurs at the speed of a normal (not accelerated) full backup. Subsequent Accelerator backups using same backup host use the track log for accelerated backup speed.

  Note: When you first enable a policy to use accelerator, the next backup (whether full or incremental) is in effect a full backup: It backs up all objects corresponding to Cloud objects queries. If that backup was scheduled as an incremental, it may not complete within the backup window.

- NetBackup retains track logs for future accelerator backups. Whenever you add a query, NetBackup does a full non-accelerated backup for the queries that are added in the list. The unchanged queries are processed as normal accelerator backups.
- If the storage unit that is associated with the policy cannot be validated when you create the policy, it is validated later when the backup job begins. If accelerator does not support the storage unit, the backup fails. In the bpbrm log, a message appears that is similar to one of the following: Storage server %s, type %s, does not support image include. Storage server type %s, does not support accelerator backup.

- Accelerator requires that the storage has the `OptimizedImage` attribute enabled.
- The Expire after copy retention can cause images to expire while the backup runs. To synthesize a new full backup, the SLP-based accelerator backup needs the previous backup.
- To detect change in metadata, NetBackup uses one or more cloud APIs per object/blob. Hence, change detection time increases with number of object/blobs to be processed. You may observe backups running longer than expected, for cases with small or no data change but having a large number of objects.
- If in your environment, for a given object, the metadata or tag are always changed (added/removed/updated) with its data. Evaluate using incremental without accelerator over incremental with accelerator from performance and cost view point.
- While creating Cloud object store policy with multiple tag-based queries, you can use few simple rules to get best effect with accelerator. Use the query builder in the policy creation page, create separate queries, one query per tag. The accelerator-based policies perform best in this configuration.

Accelerator force rescan for Cloud object store (schedule attribute)

Accelerator force rescan is a property of the full backup schedule. It is not required for Cloud object store policy.

If you use accelerator force rescan enabled full schedule with Cloud object store policy, the change detection logic considers all objects as changed. NetBackup downloads and fingerprint the data and use the track log to detect if the data is changed or un-changed. A compact backup stream that uses less network bandwidth is used between the backup host and the server.

Accelerator backup and 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. A potential catalog effect does not occur, depending on how often you use accelerator with full backups.

A full backup with accelerator completes faster than a normal full backup. It may therefore look advantageous to replace your incremental backups with accelerator full backups.

Since a full backup requires more catalog space than an incremental, replacing incremental backups with full backups increases the catalog size. When changing

your incremental backups to fulls, you must evaluate the advantage of accelerator full backups against the greater catalog space that full backups require as compared to incremental backups.

About incremental backup

NetBackup supports incremental backup for Cloud object store workloads. You can use incremental backup without enabling accelerator.

For Cloud object store workload, there are some metadata properties that do not alter the modification time for an object or blob. For example, the `Tags` in Azure blobs. Even if you change these metadata properties, the corresponding objects are not considered for the next incremental backup. This may appear like loss of data during incremental backup.

For detailed list of metadata properties that do not alter modification time for an object or blob, refer to the respective cloud provider's documentation.

For incremental backups, if an object name has path style naming scheme, then for each path, an entry is added in NetBackup. If the object, which is represented by the end node of this path style naming, has not changed since last backup (either full or last incremental, based on incremental schedule used), then that object is not included in the next incremental backup. Because of this behavior, empty paths show up in the catalog and are rendered in the browse view of restore.

About policies for Cloud object store assets

Backup policies provide the instructions that NetBackup follows to back up object(s). You can create a single policy to protect multiple buckets or containers in a Cloud object store account. You can select the objects that you want to protect using a policy. The objects are automatically discovered in the NetBackup environment and backed up. You need different policies to apply different backup logic to the objects in a Cloud object store account. Every Cloud object store account must be in at least one policy so that it can be backed up.

You can configure the following using a policy:

- Storage unit and media to use
- Backup schedules: Full, Differential incremental, and Cumulative incremental
- Backup selections: Whole bucket or container, or group of objects matching the criteria specified in the query.
  You can add the whole buckets or containers to a policy, or use queries to intelligently select the required objects inside the buckets to backup.

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 3-1 Steps for planning policies

Step	Action	Description
Step 1	Gather information about the Cloud object store account.	Gather the following information about each bucket/container: ■ The Account name: Credential and connection details mentioned in account are used to access cloud resources using REST APIs during backup. An account is associated with a single region, hence a policy can contain bucket/containers associated with that region only. ■ The bucket/container names ■ The approximate number of objects on each bucket/container to be backed up. ■ The typical size of the objects. One account may contain a large amount of data in number of objects, while the other accounts are smaller with lesser number of objects. To avoid long backup times, include the larger account in one policy and the smaller one in another policy. It may be beneficial to create more than one policy for the larger account.
Step 2	Group the objects based on backup requirements	Divide the different objects in the accounts into groups according to the different backup and archive requirements.

Table 3-1 Steps for planning policies *(continued)*

Step	Action	Description
Step 3	Consider the storage requirements	The NetBackup environment may have some special storage requirements that the backup policies must accommodate. The storage unit and volume pool settings apply to all the objects that are backed up by a policy. If objects have special storage requirements, create separate policies for the objects, even if other factors are the same, such as schedules.
Step 4	Consider the backup schedule	Create additional backup policies if the schedules in one policy do not accommodate all objects and objects in an account. Consider the following factors when deciding to create additional policies: ■ 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 objects. ■ How frequently the objects change. If some objects change more frequently than others, the difference may be enough to warrant creating another policy with a different backup frequency. ■ How long backups need to be retained. Each schedule includes a retention setting that determines how long NetBackup keeps the objects that are backed up by the schedule. Because the schedule backs up all the objects in the backup selection list, all objects should have similar retention requirements. Do not include the objects 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 backup.	You do not need to backup entire objects, unless required. Create queries to select and back up only the required object(s).

Prerequisites for Cloud object store policies

Before you begin creating a policy for a Cloud object store account, consider the following prerequisites.

- A valid Cloud object store account to access the bucket(s) and objects.
- Keep handy information about the bucket(s) and the criteria that you want to use for selecting objects from them, in the Cloud objects tab.
- You must have permission to view and select the Cloud object store account and the access host to specify the backup host for a policy, in the **Cloud objects** tab.
- Evaluate the requirement for NetBackup accelerators in your environment. If you want to use accelerators, you need to specify while creating the policy.
- If you plan to use any other backup host, other than the one used for Cloud object store account validation: Make sure that required ports are opened and configurations are done for communication from the backup host to the cloud provider endpoint using REST API calls.
- Evaluate the requirement for NetBackup multistreaming in your environment. For a given bucket, NetBackup creates one stream per query defined for the bucket in the policy. If you want to use multistreaming, you can specify while creating the policy. To use multistream, you also need to configure the number of jobs for the buckets as client in the **Client attributes** section, under primary server **Host properties**. Add the client name and set the **Maximum data streams** as required.

Creating a backup policy

Backup policies provide the instructions that NetBackup follows to back up objects. Use the following procedures to create a backup policy.

Define policy attributes like name, storage type, job priority and so on.	See "Setting up attributes" on page 29.
Schedule your backups.	See "Creating schedule attributes for policies " on page 32. See "Configuring the Start window" on page 34. See "Configuring exclude dates" on page 35. See "Configuring include dates" on page 36.

Select the account and objects to backup.	See "Configuring the Cloud objects tab" on page 36. See "Adding conditions " on page 37. See "Adding tag conditions " on page 38.

Setting up attributes

To set up attributes:

1. On the left, click **Policies**, under **Protection**.
2. Enter a name of the policy in the **Policy name** field.
3. In the **Destination** section, configure the following data storage parameters:
   - Select the **Cloud-Object-Store** option from the **Policy type** drop-down.
   - 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 the list displays **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.
     - **Any available**-If you select this option, NetBackup tries to store data on locally-attached storage units first. Select **NetBackup** or **DataStore** from the **Policy volume pool** drop-down. The **Policy volume pool** attribute specifies the default volume pool where the backups for the policy are stored. A volume pool is a set of media that is grouped for use by a single application. The volume pool is protected from access by other applications and users.

4 **Take checkpoints every**-Specify the frequency for taking checkpoints during a backup. By taking checkpoints during a backup, you can save time if the backup fails. By taking checkpoints periodically during the backup, NetBackup can retry a failed backup from the beginning of the last checkpoint. This is often quicker rather than restarting the entire job.

   The checkpoint frequency indicates how often NetBackup takes a checkpoint during a backup. The default is 15 minutes. The administrator determines checkpoint frequency on a policy-by-policy basis. When you select the checkpoint frequency, balance the loss of performance due to frequent checkpoints with the possible time lost when failed backups restart. If the frequency of checkpoints affects performance, increase the time between checkpoints.

   Checkpoints are saved at object boundaries and point to the next object in the list to be backed up. Checkpoints cannot occur in the middle of an object backup. After the object is backed up, the checkpoint is saved.

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 unchecked, 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, check 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 The **Media owner** field is available when the **Policy storage** attribute is set to **Any Available**. The **Media owner** attribute specifies which media server or server group should own the media that backup images for this policy are written to.

   - **Any**(default)-Allows NetBackup to select the media owner. NetBackup selects a media server or a server group (if one is configured).
   - **None**-Specifies that the media server that writes the image to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.

8 To activate the policy, select the option **Go into effect at**, and set the date and time of activation. The policy must be active for NetBackup to use the policy. Make sure that the date and time are set to the time that you want to resume backups.

   To deactivate a policy, clear the option. Inactive policies appear are available in the **Policies** list.

9 Select the **Allow multiple data streams** option to allow NetBackup to divide automatic backups for each query into multiple jobs. Because the jobs are in separate data streams, they can occur concurrently.

   Multistreamed jobs consist of a parent job to perform stream discovery and children jobs for each stream. Each child job displays its own job ID in the Job ID column in the **Activity monitor**. The job ID of the parent job appears in the Parent Job ID column, which is not displayed by default. Parent jobs display a dash (-) in the Schedule column.

10 Enable or disable **Disable client-side deduplication**:

   - **Enable**-The clients do not deduplicate their own data and do not send their backup data directly to the storage server. The NetBackup clients send their data to a deduplication media server. That server deduplicates the data and then sends it to the storage server.
   - **Disable**-The clients deduplicate their own data. They also send it directly to the storage server. Media server deduplication and data transport are bypassed.

11 Select the **Use accelerator** option to enable accelerator for the policy.

NetBackup accelerator increases the speed of backups. The increase in speed is made possible by change detection techniques on the client. The backup host uses the change detection techniques to identify the changes occurred between last backup and the current state of Cloud object store's object/blobs. The client sends the changed data to the media server in a more efficient backup stream. The media server combines the changed data with the rest of the client's data that is stored in previous backups.

If an object or portion of an object is already in storage and has not been changed, the media server uses the copy in storage rather than reading it from the client. The end result is a full NetBackup backup.

12 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.

Clients 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 schedule attributes for policies

This topic describes how to configure certain schedule properties for Cloud object store policies. The schedule properties vary according to your specific backup strategy and system configuration. Additional information about other schedule properties is available in the *NetBackup Administrator's Guide, Volume I.*

To create a schedule:

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 **Type of backup**:

- **Full Backup**-A complete backup of the objects that contains all of the data objects and the log(s).

- **Differential Incremental Backup**-A backup of the changed blocks since the last backup. If you configure a differential incremental backup, you must also configure a full backup.
- **Cumulative Incremental Backup**-Backs up all the changed objects since the last full backup. All objects are backed up if no previous backup was done.

4. Select the **Accelerator forced rescan** option to activate NetBackup accelerator for this policy. This option creates a checksum of the content of each object during backup. It uses the checksums for change detection. It provides a safety net by establishing a new baseline for the next accelerator backup.
5. The **Override policy storage selection** attribute works as follows:
   - **Disabled**-Instructs the schedule to use the **Policy storage** as specified on the policy **Attributes** tab.
   - **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.
6. The **Override policy volume pool** attribute works as follows:
   - **Disabled**-Instructs the schedule to override the volume pool that is specified as the Policy volume pool on the policy Attribute tab. If no policy volume pool is specified, NetBackup uses NetBackup as the default.
   - **Enabled**-Instructs the schedule to override the volume pool that is specified as the **Policy volume pool** on the policy **Attribute** tab. Select the volume pool from the list of previously configured volume pools.
7. The **Override media owner** selection attribute works as follows:
   - **Disabled**-Instructs the schedule to use the media owner that is specified as the**Media owner** in the policy **Attribute** tab.
   - **Enabled**-Instructs the schedule to override the media owner that is specified as the **Media owner** in the policy **Attribute** tab.
     Select the new media owner from the list:
     - **Any.**
       NetBackup selects the media owner, either a media server or server group.
     - **None.**
       Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.

8 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.

9 Specify a **Retention** period for the backups. This attribute specifies how long NetBackup retains the backups. To set the retention period, select a time 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.

10 The **Media multiplexing** attribute specifies the maximum number of jobs from the schedule that NetBackup can multiplex to any drive. Multiplexing sends concurrent backup jobs from one or several clients to a single drive and multiplexes the backups onto the media.

Specify a number from 1 through 32, where 1 specifies no multiplexing. Any changes take effect the next time a schedule runs.

11 Click **Add** the add the attributes or click **Add and add another** to add a different set of attributes for another schedule.

Configuring the Start window

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

For example, create different windows:

- One for the backups that open each day for a specific amount of time.
- Another for the backups that keep the window open all week.

Configuring 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 the day(s) on the 3-month calendar that you want to exclude. 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.
     - Check 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.
     - Check 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 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. Calendar-based schedules provide several run day options for determining when a task runs.

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.

Configuring the Cloud objects tab

The Cloud objects tab lets you select the Cloud object store account that you want to use to connect to cloud resources to protect object in desired buckets. NetBackup lets you make discrete selection of the buckets/containers and objects that you want to protect using the policy. You can use queries to intelligently filter and select the items that you want to protect.

NetBackup supports single backup host per policy. Hence, to distribute the load, you have to create multiple policies, and using queries you can bifurcate the load of buckets/objects being backed up across multiple backup hosts.

To configure cloud objects:

1 Select a **Cloud object store account** and **Backup host**. You can see a list of accounts and backup hosts that you are privileged to access.

2 To add buckets/containers, click **Add** on top of the **Buckets/Containers** table. In the **Add bucket/containers** dialog, do any of the following to add buckets/containers.

- To add a particular container, enter the name in the **Bucket/Container name** field, and click **Add**.
- Select one or more bucket(s)/container(s) from the **Bucket/Containers** table, and click **Add**. You can use the search box above the table to filter the list.

If the Cloud object store account credentials do not have permission to list buckets, the bucket list remains empty. But you can manually add buckets.

Note: The permissions to get list of buckets is not required to backup the objects in the bucket.

In the **Cloud objects** tab, click **Remove** in the row of any bucket/container name in the **Buckets/Containers** table to remove it from the policy. Enter a keyword in the search box to filter the table.

3 To add a query to the selected buckets/containers, click **Add query** under **Queries**.

4 Enter a name for the query, and select the buckets that you want to filter using the query.

5 In the **Select objects/blobs** table, select the option **Include all objects/blobs in the selected buckets/containers** to backup the entire bucket(s).

6 Under **Buckets with no queries**, select the buckets/containers to which you want to add queries. If a bucket is previously selected to include all queries, that bucket does not appear in this list. Click **Add condition** or **Add Tag condition** to add a condition or a tag condition. See “Adding conditions ” on page 37. and See “Adding tag conditions ” on page 38. respectively, for more details.

Adding conditions

NetBackup gives you the convenience of selectively backup the backup objects/containers inside the buckets/containers using intelligent queries. You can add conditions or tag conditions to select the objects/blobs inside a bucket/container that you want to back up.

To add a condition:

1. While creating a policy, in the **Cloud objects** tab, click **Add query**, under **Queries**.
2. In the **Add a query** dialog, enter a name for query, select the bucket(s) to which you want to apply the query. In the list of buckets, you can see only those buckets that are not selected to include all object.

   Note: While editing a query, you can see the buckets that are selected to include all objects, but the edit option is disabled.

   The **Queries** table shows the queries that you have added. You can search through the queries using values in **Query name** and **Queries** columns. The values of **Queries** column do not include the queries with **Include all objects/blobs in the selected buckets/containers** option selected.
3. Select **Include all objects in the selected buckets** option to back up all the objects in the selected bucket(s).
4. To add a condition, click **Add condition**.

   You can make conditions by using either **prefix** or **object**. You cannot use both **prefix** and **object** in the same query. Do not leave any empty fields in a condition.
5. Select **prefix** or **object** from the drop-down, enter a value in the text field. Click **Condition** to add another condition. You can join the conditions by the boolean operator OR.
6. Click **Add** to save the condition.

Adding tag conditions

You can add tag conditions to select the object/blob that you want to back up, using key-value pairs, and boolean conditions.

To add tag conditions:

1. While creating a policy, in the **Cloud objects** tab, click **Add query**, under **Queries**.
2. In the **Add a query** dialog, enter a name for query, select the bucket(s) to which you want to apply the query. In the list of buckets, you can see only those buckets that are not selected to include all object.
3. Select **Include all objects in the selected buckets** option to back up all the objects in the selected bucket(s).

4. To add a tag condition, click **Add Tag Condition**.
5. Enter values for **Tag Key** and **Tag Value** to create the condition. The boolean operator **AND** joins the values. NetBackup backs up objects with the matching key-value pairs.
6. Click **Tag condition** to add more conditions. You can use the boolean **AND** or **OR** parameters to connect the tag conditions.
7. Click **Add** to save the condition.

Example of conditions and tag conditions

Here is an example to illustrate the use of conditions and tag conditions.

Consider the container/bucket has the following files/directories:

- Following blobs tagged with "Project": "HR" tag
  - OrganizationData/Hr/resumes/resume1_selected.pdf
  - OrganizationData/Hr/resumes/resume2_rejected.pdf
  - OrganizationData/Hr/resumes/resume3_noupdate.pdf
- Following blobs tagged with "Project": "Finance" tag value
  - OrganizationData/Fin/accounts/account1/records1.txt
  - OrganizationData/Fin/accounts/account2/records2.txt
  - OrganizationData/Fin/accounts/account3/records3.txt
  - OrganizationData/Fin/accounts/monthly_expenses/Jul2022.rec
  - OrganizationData/Fin/accounts/monthly_expenses/Aug2022.rec
- Following blobs tagged with "Project": "Security"
  - The blob Getepass.pdf: Has one more tag with “TypeOfData”:”ID_Cards” so this is tagged with two tags (that is: Security and ID_Cards)
  - OrganizationData/newJoinees/tempPassesList.xls
- Following blobs tagged with "Project": "Environment"
  - EnvironmentContribution.xls
  - NewPlantedTrees.xls

Example prefix conditions:

- Case 1: To backup all resumes irrespective of their status (like, selected or rejected) from `OrganizationData` add the query:

`prefix Equal to OrganizationData/Hr/resumes/resume`

Result: All records that start with `OrganizationData/Hr/resumes/resume` are backed up.

- Case 2: To backup all resumes and records from Fin and HR, add any of the following queries:

  `prefix Equal to OrganizationData/Hr/resumes/resume`

  Or

  `prefix Equal to OrganizationData/Fin/accounts/account1/rec`

  Note: You can add multiple prefixes with OR conditions.

Result: All records starting with `OrganizationData/Hr/resumes/resume` or `OrganizationData/Fin/accounts/account1/rec` are backed up.

Example object conditions:

To backup a specific object/blob add the following query:

```
object Equal to
OrganizationData/Fin/accounts/monthly_expenses/Jul2022.rec
```

Result: Only the blob with the name `Jul2022.rec` is selected.

Example tag conditions:

- Case 1: To backup all blobs tagged with "Project": "Finance", add the following query:

  tagKey Equal to 'Project' and tagVal Equal to 'Finance'

  Result: All object/blobs tagged with "Project" = "Finance" are selected.

- Case 2: To backup data that matches with project Finance or Security, add the query:

  tagKey Equal to 'Project' and tagValue eq 'Finance' **OR** tagKey Equal to 'Project' and tagValue eq 'Security'

  Result: All object/blobs tagged with "Project": "Finance" or "Project":"Security" are selected.

- Case 3: To backup data from “Project”:”Security” and “TypeOfData”:”ID_Cards” add the queries:

  (tagKey Equal to 'Project' and tagValue Equal to 'Security') **AND** (tagKey Equal to 'TypeOfData' and tagValue Equal to 'ID_Cards')

  Result: Data with tag “Project”:” Security” and “TypeOfData”: “ID_Cards” are selected.

Managing Cloud object store policies

You can add, edit, delete, copy, and deactivate policies. You can also perform manual backup for a policy.

View Cloud object store policies

1. On the left, click **Policies**. All the policies that you have privileges to view are displayed.
2. To filter the table for Cloud object store policies, click the filter icon and select **Cloud-Object-Store**.

   Use the search box at the top of the table to search for a policy.

To edit a Cloud object store policy, select the policy. Then click **Edit**.

See "Creating a backup policy" on page 28.

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 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 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 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 manual backup:

1. On the left, click **Policies**. All the policies that you have 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, 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.

Chapter 4

Recovering Cloud object store assets

This chapter includes the following topics:

Prerequisites for recovering Cloud object store objects

Ensure that the following prerequisites are satisfied, before you start recovery.

- Keep handy information about the destination bucket(s) or container(s) that you want to use for recovery.
- Decide on your object selections for recovery. You can recover objects by selecting all objects or blobs from selected image. Alternatively, select individual objects, select all objects under set of folder(s), or all objects matching set of prefix(es).
- A valid Cloud object store account to access the bucket(s) or container(s) and objects/blobs. You can add the Cloud object store account related information to NetBackup while creating the account. The permission required for restoring differs from ones required for backup, if it helps you can create separate Cloud object store account for recovery.
- Ensure that you have permission to view and select the Cloud object store account and the access host. To be able to select recovery host for a policy, in the **Cloud objects** tab.
- If required, you can use a different recovery host than the one used for Cloud object store account validation. Ensure that the new recovery host has the

required ports opened and configured for communication from the backup host to the cloud provider endpoint, using REST API calls.

- You can plan to start multiple restore jobs in parallel for better throughput. You can select objects for recovery as individual objects, or using folder or prefix.

Recovering Cloud object store assets

You can recover Cloud object store assets to the original or a different bucket or container. You can also restore each of the objects to different buckets or containers.

To recover assets:

1. On the left, click **Recovery**. Under **Regular recovery**, click **Start recovery**.
2. In the Basic properties page, select **Policy type** as **Cloud-Object-Store**.
3. Click the **Buckets/Containers** field to select assets to restore.
   - In the Add bucket/container dialog, the default option, displays all available bucket/containers for recovery. You can search the table using the search box.
   - To add a specific bucket or container, select **Add the bucket/container details** option. Enter the name of the bucket/container, select the cloud provider, and enter the Cloud object account name.

     Note: In a rare scenario, if you cannot find the required bucket listed in the table for selection. But you can see the same bucket listed in catalog view as backup ID. You can select the bucket by manually entering bucket name, provider ID, and the Cloud object store account name as per the backup ID. The backup ID is formed as
     `<providerId>_<cloudAccountname>_<BucketName>_<timestamp>`
4. Click **Add**, and then, click **Next**.
5. In the Add objects page, select the **Start date** and the **End date** of the period from which you want to restore.

   (Optionally) Enter a keyword phrase to filter the images, and click **Apply**.
6. Click **Backup history**, and select the required images for recovery from the **Backup history** dialog. Click **Select**.
7. (Optional) Click **Add objects and folders**, and select the required objects to recover from the **Add Object/blobs and folders** dialog. Select **Include all objects/blobs and folders** to include all available assets. You can use the left navigation tree structure to filter the table. Click **Add**.

8. (Optionally) Click **Add** prefix. In the Add prefix dialog, enter a prefix in the search box to display relevant results in the table. Click **Add**, to select all the matching prefixes displayed in the table for recovery. The selected prefixes are displayed in a table below the selected objects/blobs. Click **Next**.
9. In the Recovery options page, you can select whether you want to restore to the source bucket of container or user different ones. These are the **Object restore options**:
   - **Restore to the original bucket or container**: Select to recover to the same bucket or container from where the backup was taken. Optionally, add a prefix for the recovered assets in the **Add a prefix** field.
   - **Restore to a different bucket or container**: Select to recover to a different bucket or container than the one from where the backup was taken.
     - You can select a different **Cloud object store account** as destination, from the list above.
     - Select a destination **Bucket/Container name**. You can use different Cloud object store accounts that can access the original bucket. This method also helps you to make accounts with limited and specific permissions for backup and restore. In this case, you can provide the same bucket as original to restore to original bucket/container.
     - Optionally, add a prefix for the recovered assets in the **Add a prefix** field.
   - **Restore object/blobs or prefixes to different destinations**: Select to recover each of your selected assets to different destinations.
     - You can select a different **Cloud object store account** as destination, from the list above.
     - Click **Edit object destination**, enter the **Destination** and **Destination bucket/container** name. Click **Save**.

   Note: If you have selected **Include all objects/blobs and folders**, in step 7, the **Restore objects/blobs or prefixes to different destinations** option is disabled.
10. Select a **Recovery host**.
11. Optionally, to overwrite any preexisting object or blobs using the recovered assets, select **Overwrite existing objects/blobs**.

12 (Optional) To override the default priority of the restore job, select **Override default priority**, and assign the required value.

13 In the Review page, check all the summary of all the selections that you made, click **Start recovery**.

You can see the progress of the restore job in the Activity monitor.

Chapter 5

Troubleshooting

This chapter includes the following topics:

Recovery for Cloud object store using web UI for original bucket recovery option starts but job fails with error 3601

Explanation

Can be because any of the four reasons: Happens as the , or

- Cloud object store account required to connect to cloud to do recovery does not exist.
- The Cloud object store account used during backup of the bucket does not exist in NetBackup domain.
- This is the target domain with AIR configuration or DR scenario.
- The Cloud object store account was deleted.

Workaround

Create the Cloud object store account with same name and provider as the original Cloud object store account and retry the recovery.

Recovery Job does not start

Explanation

Recovery to the original bucket gives the error "Unable to retrieve asset details". Even if the Cloud object store account with same name is used during backup.

Workaround

Do the following:

1. Use the same Cloud object store account in web UI.
2. Try to recover to a different bucket in the same account. This action refreshes the cache.

   You can force cache refresh even by using asset API with no-cache to fetch all assets of cloudObjectStoreAccount (/netbackup/asset-service/workloads/cloud-object-store/assets/?filter=assetType eq 'cloudObjectStoreAccount') Make sure that the account is listed in output.
3. Now, again use the original bucket recovery option and perform the recovery.

Restore fails: "Error bpbrm (PID=3899) client restore EXIT STATUS 40: network connection broken"

Explanation

Delay in accessing the backup image and uploading the blobs for the restore. In this process, the bptm is getting timed out.

Workaround

If restore fails with network error message in the Activity monitor, change system configuration time-out to 900/1200 or appropriate high value, in seconds, and start a new restore job. Steps to set time-out are as following: **Menu** > **Hosts** > **Host properties** > Select **Primary/ Media server** > select the **Timeout** option > set the time out value -> click **Save**. Refer to the *Web UI Administrator Guide* for details.

Access tier property not restored after overwrite existing to original location

Explanation

Object with access tier Cool, overwritten by a restore with Hot access tier, does not change the access tier to Hot, it remains Cool.

Workaround

In case of Azure cloud storage, when we have an object with accessTier as cool, and we try to upload object/blob of the same name with Hot(inferred) accessTier with overwrite option, accessTier remains as cold. New access tier does not get set. This behavior is observed when file is uploaded from portal. It does not change accessTier from cool to Hot (inferred) when the Overwrite option is selected on Azure portal.

Reduced accelerator optimization in Azure for OR query with multiple tags

Explanation

When Cloud object store policy has at least one query having multiple tag conditions combined using "OR" operator, the backups of Azure containers using accelerator enabled Cloud object store policy, shows loss of acceleration or backs up unchanged data.

This happens as the ordering of objects across multiple tags is not as expected for accelerator. Few objects that are not found in the tracklog even if they exist in tracklog and hence backed up repeatedly without getting accelerator benefit for these objects.

Workaround

Do not use OR condition while combining multiple tag conditions for Azure. Instead create separate query per tag.

For example,

Say that you have the following query (tagKey eq 'type' and tagValue eq 'text') or (tagKey eq 'type' and tagValue eq 'none') with say queryname datatype

You can create two queries say by name datatype-text with query (tagKey eq 'type' and tagValue eq 'text') and datatype-none with query (tagKey eq 'type' and tagValue eq 'none')

Note: This results in first backup which is without any acceleration for these new queries. For subsequent backups you can see the problem is resolved.

Backup is failed and shows a certificate error with Amazon S3 bucket names containing dots (.)

Workaround

Use any of these two workarounds:

- `Use path style URL to access bucket`: Since path style URL adds bucket as part of URL path and not as hostname, we did not get any SSL issues even for buckets with a . (dot) in the name. However, NetBackup default configuration uses Virtual style for all dual stack URLs like `s3.dualstack.<region-id>.amazonaws.com`. We can add older s3 URL as path style and can connect with bucket with a (.) in the name. To do this we can add region with plain s3 endpoint (`s3.<region-id>.amazonaws.com`) and selecting URL Access Style as path style.
- Disable SSL: This workaround is not the recommended one since it replaces the secure endpoint with unsecure/unencrypted endpoint. After turning off SSL it disables peer host validation of server certificate. It bypasses the hostname match for virtual host style URL of bucket (bucket.123.s3.dualstack.us-east-1.amazonaws.com) with subject name in certificate (*. s3.dualstack.us-east-1.amazonaws.com).

Azure backup job fails when space is provided in tag query for either tag key name or value.

Workaround

Do not use spaces in tag query for either tag key name or value, for Azure backup jobs.

The Cloud object store account has encountered an error

Explanation

In web UI, the Cloud object store account status is shown as: The Cloud object store account has encountered an error, see user documentation, and re-create the account.

You cannot edit the Cloud object store account in this state. All jobs corresponding to the Cloud object store account keep failing.

Cause

The Cloud object store account goes to error state when:

- The alias corresponding to Cloud object store account is accidentally deleted using csconfig CLI.
- The alias corresponding to the Cloud object store account is accidentally updated using csconfig CLI.

Note: It is recommended not to use csconfig CLI to update alias corresponding to Cloud object store account. Correct way to update the same is through the Edit workflow or create-or-update API. Alias with same name as Cloud object store account is the alias corresponding to Cloud object store account.

Workaround

The NetBackup domain name must be unique across Cloud object store account, Cloud storage server, or MSDP-C LSU. They share single namepace. Hence we can have following usage scenarios:

`Case 1`: When there is no valid Cloud storage server or MSDP-C LSU with same name as Cloud object store account in the environment.

- Gather the Cloud object store account details as per your environment and cross-check the details obtained.

- Optionally, if the Alias corresponding to the Cloud object store account exists, use csconfig CLI and note down details of alias.
  - Use following command to list all instances for the type and locate the Cloud object store account and its instance:
    `<install-path>/csconfig cldinstance -i -pt <provider_type>`
  - Use following command to get the details of instance and the Cloud object store account:
    `<install-path>/csconfig cldinstance -i -in <instance name>`
  - Validate the details with the gathered information.
  - Delete the Alias using following command:
    `<install-path>/cscpnfig cldinstance -at <api_type> -rs -in <instance_name> -sts <cloud_object_store_account_name>`
- Delete the Cloud object store account which is in error state.
- Create the Cloud object store account using the noted details.

`Case 2`: When you have valid and in use Cloud storage server or MSDP-C LSU with same name as Cloud object store account in the environment.

- You cannot re-use the same name.
- You need to gather the Cloud object store account details as per your environment.
- Identify the new name for the Cloud object store account.
- Delete the Cloud object store account which is in error state. Remove the account from the policy.
- Create the Cloud object store account using new name and details gathered. Assign this account to the same policy that the old account used.
- This changes the Client Name used for the bucket starting from the next backup onwards.
- NetBackup identifies the old backups using the old account name.

Bucket list empty when selecting it in policy selection

Explanation:

When you configure a Cloud object store account by adding a region entry, without specifying a correct region location constraint. The account gets added successfully, because some private cloud might not have region configured.

When you are using such an invalid region in an account, the List bucket may return empty.

Workaround:

Do the following:

1. Call the `getBucketLocation` API on the bucket, to retrieve the correct location constraint for your account configuration.

   If the API returns a blank location constraint, use 'us-east-1' as region location constraint.

2. Correct the region details by editing the account configuration. See “Adding Cloud object store accounts” on page 10.

3. To edit cloud configuration, do the following:

   - On the left, click **Host Properties**.
   - Select the required primary server and connect it. Click **Edit primary server**.
   - Click **Cloud storage**.
   - Optionally, enter your cloud provider name in the search field, to filter the list.
   - In the row corresponding to your cloud provider service host enter the correct region details and save.

Alternatively, delete the account and recreate it with the correct region location constraint.

Creating second account on Cloudian fails by selecting existing region

Explanation:

After adding a Cloud object storage account for Cloudian by adding region with us-east-1 location constraint. If you try to reuse the same region and create a second account, account creation fails.

This happens because region listing API is converting region's location constraint 'us-east-1' to "<blank> while showing in web UI. You can see added region location constraint was us-east-1 and one which is listed has blank location constraint field. Account created by selecting such region from list fails.

Workaround:

Use the NetBackup Asset Query APIs to create account. Example region details part which can be provided in payload:

```
"s3RegionDetails": [
  { "regionId": "us-east-1",
    "regionName": "<region name same as listed from prior account>",

   "serviceHost": "<service host same as listed from prior account>"

  }
]
```

You can obtain API DOC from schema API:

```
https://<primary-server-hostname>/netbackup/asset-service
/workloads/saas/schemas/create-or-update-assets-named-query-request
```

Restore failed with 2825 incomplete restore operation

Not all objects restored from backup image. Restore failed with 2825 incomplete restore operation.

Explanation:

This error can occur due to multiple reasons. Most likely scenario for this error is when a cloud API initiated by NetBackup during restore returns an error like HTTP 400 status code (Bad Request). The reasons can vary with each cloud vendor. For example, GCP supports different Content-Language metadata as compared to AWS. In some cases, the error can also occur depending on the features enabled or disabled on a specific cloud account or bucket.

`nbcosp` logs show the following messages:

```
{"level":"warn","error":"InvalidArgument: Invalid argument.\n\tstatus
 code: 400, request id: , host id: ","object
key":"meta-user-defined/t2.rtf","time"…}

nbtar logs will have below type of errors-
15:56:15.739 [22496.22496] <16> operation_to_cloud_by_type: ocsd
reply with error, error_code: 400
15:56:15.739 [22496.22496] <16> CloudObjectStore::InitMultiPartUpload:
 operation_to_cloud_by_type() failed, status=3600
15:56:15.739 [22496.22496] <16> CloudObjectStore::ObjectOpen:
InitMultiPart Upload call failed with status = 3600
15:56:15.739 [22496.22496] <16> cCloudApiRestoreHandler::writeOpen:
 ERR - ObjectOpen failed with error code [3600]
```

Workaround:

When the error is not fatal, the restore job is a partial success. Check the Activity Monitor to see the list of objects that cannot be restored. Try restoring to a different location (bucket/container or different account) to check if the problem is with the destination cloud account or bucket setting.

When the error is fatal, the restore job fails. Check the `nbcosp` logs to determine the object for which the restore has failed. Use granular object selection for the next restore and skip the earlier failed object while selecting the objects.

Refer to your cloud provider documentation to check if you use a feature or any metadata that the cloud vendor does not support completely, or if it needs any more configuration. Fix the object with right attributes in the Cloud object store and start a new backup job. Once this backup completes, the objects can be restored without this workaround.

Bucket listing of cloud provider fails when adding bucket in Cloud objects tab

Explanation

The most common reason for failure in bucket listing is when cloud credentials provided to NetBackup do not have permission to list buckets.

Another reason is when the cloud provider does not support proper DNS entries for endpoints. Similarly, a wrongly configured DNS or even a virtual-hosted style naming implying that no request can be issued to the cloud provider without providing a bucket name as host name. An example of such a cloud endpoint is:

`s3-fips.us-east-1.amazonaws.com`

Workaround

Although the bucket list is not available, you can always manually add buckets in the Cloud objects tab for backup.

When it is a DNS issue, you can optionally list buckets using a temporary workaround by adding IP hostname-mapping entry in the `/etc/hosts` file. When only virtual-hosted style requests are supported, first prefix the endpoint using a random bucket name, when using commands like ping, dig, nslookup to determine the IP of the cloud endpoint. For example,

```
ping randombucketname.s3-fips.us-east-1.amazonaws.com
```

You can then add the resulting IP along with the actual endpoint name (without the random bucket name prefix) in `/etc/hosts` file.

Note that this is a temporary workaround to edit DNS entries on the computer for bucket listing. Remove them after the policy configuration is done, unless the cloud endpoint is a private cloud setup that can use static IP addresses permanently.

AIR import image restore fails on the target domain if the Cloud store account is not added in target domain.

Error

Cannot perform the Cloud object store protection (COSP) operation, skipping the object:[<object name>], error: [3605]

Explanation

Cloud object store account is not present in target domain with same name as in source domain.

Workaround

Solution 1:

Create the Cloud object store account in target domain with the same name as in the source domain and perform restore. See “Adding Cloud object store accounts” on page 10.

Solution 2:

If you have a Cloud object store account with valid credentials on the target domain, do the following:

1. In the **Recover** tab, select the Bucket/Container with source account name. Click **Next**.
2. Select the backup image and add objects and folders, or prefix. Click **Next**.
3. On the **Recovery options** page, select the option: **Restore to a different bucket or container**. Select another existing Cloud object store account for restore.