

NetBackup™ for Cloud Object Store Administrator's Guide

Release 11.2

NetBackup™ for Cloud Object Store Administrator's Guide

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	Introduction	8
	Overview of NetBackup protection for Cloud object store	8
	Features of NetBackup Cloud object store workload support	9
Chapter 2	Managing Cloud object store assets	13
	Planning NetBackup protection for Cloud object store assets	14
	Enhanced backup performance in 11.0 or later	14
	Prerequisites for adding Cloud object store accounts	15
	Configuring buffer size for backups	16
	Configure a temporary staging location	17
	Configuring advanced dynamic multi-streaming parameters	19
	Permissions required for Amazon S3 cloud provider user	21
	Permissions required for Azure blob storage	22
	Permission required for Azure Data Lake Storage	23
	Permissions required for GCP	26
	Limitations and considerations	27
	Adding Cloud object store accounts	29
	Creating cross-account access in AWS	34
	Check certificate for revocation	35
	Managing Certification Authorities (CA) for NetBackup Cloud	35
	Adding a new region	38
	Manage Cloud object store accounts	38
	Scan for malware	40
	Backup images	40
	Assets by policy type	41
Chapter 3	Protecting Cloud object store assets	44
	About accelerator support	45
	How NetBackup accelerator works with Cloud object store	45
	Accelerator notes and requirements	46
	Accelerator force rescan for Cloud object store (schedule attribute)	47
	Accelerator backup and NetBackup catalog	48

Calculate the NetBackup accelerator track log size	48
About incremental backup	49
About dynamic multi-streaming	49
Enable auto-resume for incomplete backup jobs	50
About object change tracking	51
Configuring object change tracking	52
Configuring access permissions for the buckets	52
Configuring access policy on the log bucket	52
Configuration guidelines for IBM Storage Ceph	53
Enable bucket logging for source buckets	54
Creating policy for the log bucket	56
Additional storage requirements at the staging location	58
Configuring bucket logging in IBM Storage Ceph	58
Maintaining the log bucket	59
Configuring NetBackup for object change tracking	59
Configuring NetBackup policy for object change tracking	60
Verifying object change tracking in the Activity monitor	60
Scenarios when NetBackup overrides object change tracking	60
About storage lifecycle policies	61
Adding an SLP	61
About policies for Cloud object store assets	64
Planning for policies	64
Prerequisites for Cloud object store policies	67
Creating a backup policy	68
Policy attributes	68
Creating schedule attributes for policies	72
Configuring the Start window	74
Adding, changing, or deleting a time window in a policy schedule	74
Example of schedule duration	76
Configuring the exclude dates	76
Configuring the include dates	77
Configuring the Cloud objects tab	78
Adding conditions	80
Adding tag conditions	81
Examples of conditions and tag conditions	82
Managing Cloud object store policies	84
Copy a policy	84
Deactivating or deleting a policy	84
Manually backup assets	85

Chapter 4	Recovering Cloud object store assets	86
	Prerequisites for recovering Cloud object store objects	86
	Configuring Cloud object retention properties	87
	Recovering Cloud object store assets	87
Chapter 5	Troubleshooting	92
	Credential validation failed. Provide correct Cloud account name and credentials.	94
	Error 5549: Cannot validate bucket logging information	95
	Error 5576: The maximum number of concurrent jobs specified for a storage unit, must be greater than or equal to the number of streams specified in the policy.	95
	Error 5579: Falling back to object listing for change detection, not considering object change tracking for this bucket, specified in the policy.	96
	Error 5580: The specified failover strategy for the storage unit group is incompatible with the Cloud object store policy, with dynamic multi-streaming.	96
	Error 5545: Backup failed as NetBackup cannot parse records from the log object	97
	Error 5541: Cannot take backup, the specified staging location does not have enough space	97
	Error 5537: Backup failed: Incorrect read/write permissions are specified for the download staging path.	98
	Error 5538: Cannot perform backup. Incorrect ownership is specified for the download staging path.	98
	Reduced acceleration during the first full backup, after upgrade to versions 10.5 and 11.	98
	After backup, some files in the <code>shm</code> folder and shared memory are not cleaned up.	99
	After an upgrade to NetBackup version 10.5, copying, activating, and deactivating policies may fail for older policies	99
	Backup fails with default number of streams with the error: Failed to start NetBackup COSP process.	100
	Backup fails, after you select a scale out server or Snapshot Manager as a backup host	100
	Backup fails or becomes partially successful on GCP storage for objects with content encoded as GZIP.	100
	Recovery for the original bucket recovery option starts, but the job fails with error 3601	101
	Recovery Job does not start	101

Restore fails: "Error bpbrm (PID=3899) client restore EXIT STATUS 40: network connection broken"	102
Access tier property not restored after overwriting the existing object in the original location	102
Reduced accelerator optimization in Azure for OR query with multiple tags	103
Backup failed and shows a certificate error with Amazon S3 bucket names containing dots (.)	103
Azure backup jobs fail when space is provided in a tag query for either tag key name or value.	104
The Cloud object store account has encountered an error	104
The bucket is list empty during policy selection	106
Creating a second account on Cloudian fails by selecting an existing region	107
Restore failed with 2825 incomplete restore operation	107
Bucket listing of a cloud provider fails when adding a bucket in the Cloud objects tab	108
AIR import image restore fails on the target domain if the Cloud store account is not added to the target domain	109
Backup for Azure Data Lake fails when a back-level media server is used with backup host or storage server version 10.3	110
Backup fails partially in Azure Data Lake: "Error nbpem (pid=16018) backup of client	110
Recovery for Azure Data Lake fails: "This operation is not permitted as the path is too deep"	111
Empty directories are not backed up in Azure Data Lake	111
Recovery error: "Invalid alternate directory location. You must specify a string with length less than 1025 valid characters"	111
Recovery error: "Invalid parameter specified"	112
Restore fails: "Cannot perform the COSP operation, skipping the object: [/testdata/FxtZMidEdTK]"	112
Cloud store account creation fails with incorrect credentials	113
Discovery failures due to improper permissions	114
Restore failures due to object lock	114

Introduction

This chapter includes the following topics:

- [Overview of NetBackup protection for Cloud object store](#)
- [Features of NetBackup Cloud object store workload support](#)

Overview of NetBackup protection for Cloud object store

NetBackup enables backup and restore of cloud object stores on both on-premises and cloud-based object store services. You can deploy NetBackup in the same cloud network as the object store or provide HTTP(S) connectivity between the object store service endpoint and the backup host or scale-out server. You can also deploy NetBackup outside the vendor's cloud.

Note: Cloud vendors may charge high data egress fees for transferring data out of their network. Before configuring a backup policy that moves data to another cloud, cloud region, or on-premises data center, check your cloud provider's data-out pricing.

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 stores, IBM Storage Ceph, and so on. For a complete list of compatible object stores, refer to the NetBackup Hardware Compatibility List (HCL).

The protected objects in Azure Data Lake are referred to as files and directories, even though the underlying object is of type blob.

Features of NetBackup Cloud object store workload support

Table 1-1 Salient features

Feature	Description
Integration with NetBackup's 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. You do not need to be a NetBackup administrator to manage Cloud object stores.
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 on security. For protecting a single Azure Blob Storage account, the storage account and access key must be specified. To protect the Azure blob storage account, the supported authentication mechanisms are Access key, Service Principal, and Managed Identity. For all S3 API-compliant cloud vendors, the Access key and Secret Key are supported. For Amazon S3, the Access Key, IAM role, and Assume role (for cross-AWS account) mechanisms of authentication are supported. For a complete list, see the NetBackup Compatibility Lists.
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 back up all the objects in them. You can also configure some buckets and containers with intelligent queries to identify objects based on: - Object name prefix - Entire object name - Object tags - Files and directories in Azure Data Lake

Table 1-1 Salient features *(continued)*

Feature	Description
Fast and optimized backups	In addition to full backup, NetBackup also supports different types of incremental schedules for faster backups. Accelerator feature is also supported for the Cloud object store policies. Enable checkpoint restart in the policy to be able to resume a failed or suspended job from the last checkpoint.
Granular restore	NetBackup makes it easy to restore all objects in a bucket or container. It also lets you select which objects to restore by using a prefix, folder, or object-based views. You can narrow down a selection of backup images for restoration in NetBackup by providing a date and time range.
Restore options	NetBackup restore the object store data along with their metadata, properties, tags, ACLs, and object lock properties. NetBackup supports adding an arbitrary prefix to all objects when restoring. Consequently, it restores objects with a distinct name when it is desired to avoid any interference with the original objects. The Azure Data Lake files and directories, however, do not require a prefix. Instead, the files and directories are restored to a specified alternate location. By default, NetBackup skips overwriting objects that already exist in the Cloud object store to conserve bandwidth and cloud costs. You can modify this default behavior by using the Overwrite option, thereby enabling the restoration of copies to overwrite the copies stored in the Cloud object store.
Alternate location restores	You can select restore objects to: ■ To the same bucket or container ■ To a different bucket or container in the same account or subscription ■ To a different bucket or container in a different account or subscription.

Table 1-1 Salient features *(continued)*

Feature	Description
Support for malware scan before recovery	You can run malware scan of the selected files or folders for recovery as part of the recovery flow from Web UI and decide the recovery actions based on malware scan results.
Dynamic multi-streaming	This feature allows multiple backup streams to operate simultaneously for a single object store bucket/container, optimizing backup performance to meet the backup windows. Here are some advantages of Dynamic multi-streaming: ■ Implicitly distributes the objects for backup across multiple streams. ■ Automate stream creation and data distribution across streams. ■ Prevents stream starvation.
Scalability support for the backup host	NetBackup Cloud object store protection supports configuring the NetBackup Snapshot Manager as a scalable backup host for cloud deployments, along with the media server. If you have an existing NetBackup Snapshot Manager deployment in your environment, you can use that as a backup host for Cloud object store policies. With NetBackup Snapshot Manager as the backup host, you do not need to configure multiple backup hosts for large jobs or create multiple policies to distribute the load across these backup hosts. Snapshot Manager can increase the number of data mover containers during a backup operation, and then reduce them when the protection tasks are completed.
Object lock	This feature lets you retain the original object lock properties and also provides an option to customize the object lock properties. If you use object lock properties on the restored objects, you can't delete those objects until the retention period is over, or the legal holds are removed. You can use the Object lock and retention properties without any configuration during policy creation and backup.

Table 1-1 Salient features *(continued)*

Feature	Description
Quick object change scan	This feature significantly speeds up the NetBackup Accelerator resulting in faster backups. Apart from an object's modification time, NetBackup also checks for changes in the object's tags and user attributes to determine whether the object qualifies for an incremental backup. In those application environments, where metadata changes do not occur after the objects are created, you can use this option to skip these metadata checks that significantly increases the backup speed. Enabling this option speeds up change identification for Accelerator, as it skips the tags and attribute checks of the objects since the previous backup. However, the skipped objects are not backed up. Those object metadata changes which does not result in change in the modification time of the object are skipped. For more information, see the Knowledge article: https://support.cohesity.com/s/article/article-100073853.html
Backup host	A backup host is used for credential validation, backup, and recovery of the Cloud object store assets. You can configure a NetBackup media server as a backup host. NetBackup's object store agent resides on this backup host and interacts with various Cloud object store vendors and services. Starting NetBackup 11.0.0.1, you can configure a RHEL or SUSE media server as a backup host.
Object Change Tracking	Object store vendors provide mechanisms to track changes in buckets or containers for use cases such as auditing and inventory. NetBackup leverages these mechanisms to efficiently identify changed objects without scanning the entire bucket, significantly improving backup performance.

Managing Cloud object store assets

This chapter includes the following topics:

- [Planning NetBackup protection for Cloud object store assets](#)
- [Enhanced backup performance in 11.0 or later](#)
- [Prerequisites for adding Cloud object store accounts](#)
- [Configuring buffer size for backups](#)
- [Configure a temporary staging location](#)
- [Configuring advanced dynamic multi-streaming parameters](#)
- [Permissions required for Amazon S3 cloud provider user](#)
- [Permissions required for Azure blob storage](#)
- [Permission required for Azure Data Lake Storage](#)
- [Permissions required for GCP](#)
- [Limitations and considerations](#)
- [Adding Cloud object store accounts](#)
- [Manage Cloud object store accounts](#)
- [Scan for malware](#)

Planning NetBackup protection for Cloud object store assets

This section elaborates the tasks that you need to perform to deploy NetBackup to protect your Cloud object store assets.

Table 2-1 Steps for NetBackup deployment

Step	Action	Description
Step 1	Verify the operating system and platform compatibility.	See the NetBackup Compatibility Lists .
Step 2	Install NetBackup	See the <i>NetBackup™ Installation Guide</i>
Step 3	Configure the required permissions and credentials.	See “ Prerequisites for adding Cloud object store accounts ” on page 15.
Step 4	Identify the buckets and containers that you want to protect.	Make a list of the buckets and containers that you want to protect with NetBackup, and include them in the Cloud object store accounts that you create in Step 5.
Step 5	Create Cloud object store accounts.	See “ Adding Cloud object store accounts ” on page 29.
Step 6	Create policy.	See “ Creating a backup policy ” on page 68.

Enhanced backup performance in 11.0 or later

NetBackup 11.0 introduces significantly improved Cloud object store backups by decentralizing object handling and backups.

The new design facilitates parallel processing and object copy, which significantly reduces latency, increases throughput, and improves overall backup efficiency.

NetBackup uses two distinct modules that independently exchange Cloud object store data and metadata over a staging location on the backup host. These two modules can fetch objects, manage the staging area, incorporate accelerator backup logic, and back up data to target, all independent of each other.

You can now skip tracking object metadata changes. Apart from an object's modification time, NetBackup also checks for changes in the object's tags and user attributes to determine whether the object qualifies for an incremental backup. In those organizations, where metadata changes do not occur after the objects are created, you can use this feature to skip the metadata checks that significantly

increases the backup speed. You can configure this setting while creating policies for Cloud object store. See [“Policy attributes”](#) on page 68.

You must configure a temporary staging location to use this feature. See [“Configure a temporary staging location”](#) on page 17. Step 7.

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.
Here, the service host is the host name of the Cloud object storage API endpoint that is provided by the cloud provider. For example, in the AWS public S3 endpoint URL: `https://s3.us-east-1.amazonaws.com`, the part: `s3.us-east-1.amazonaws.com` is the service host.
For a private cloud setup, the URL may be like:
`https://s3.us-east-1.amazonaws.com/tenant123/`. Here the service host is: `s3.us-east-1.amazonaws.com/tenant123/`.
- Note 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
 - Assume role EC2: For Amazon and Amazon Gov
 - Credential Broker: For Amazon Gov
 - Service Principal: For Azure
 - Managed Identity: For Azure
- If you plan to use a proxy for communication with cloud endpoints, gather the required details of the proxy server.
- Get the Cloud account credentials, and any additional required parameters, as per the authentication type. These credential details should have the required permissions recommended in NetBackup documentation.
See [“Permissions required for Amazon S3 cloud provider user”](#) on page 21.
See [“Permissions required for Azure blob storage”](#) on page 22.
See [“Permissions required for GCP”](#) on page 26.

- Make sure that the required outbound ports are open, and configurations are done for communication from the backup host or scale-out server to the cloud provider endpoint using REST API calls.
 - On the backup host, S3 or Azure storage URL endpoints use the HTTPS default port 443. For a private cloud provider, this port can be any custom port that is configured in the private cloud storage.
 - If you use a proxy server to connect to the cloud storage, you need to allow that port. You can provide the proxy server-related details in NetBackup, while creating a Cloud object store account.
 - The certificate revocation status check option uses the OCSP protocol, which typically uses HTTP port 80. Ensure that the OCSP URL is reachable from the backup host.

Configuring buffer size for backups

The read or write buffer size for object stores influence the performance of NetBackup in the cloud object stores.

Note: The configurations described in this section are not required if your backup host is NetBackup version 11.0 or higher, and your policy does not have Dynamic multi-streaming enabled. For policies with Dynamic multi-streaming, see the section *Configuring a temporary staging location*.

By default, NetBackup creates buffers of 4 MB. You can use this default buffer size, if most of the objects in your buckets/containers are less than 4 MB. If your buckets/containers have a large number of objects greater than 4 MB then you can increase the buffer size up to 64 MB.

To configure buffer size:

- 1 Open the `/usr/openv/netbackup/bp.conf` file on the backup host.

You can identify the backup host from the corresponding policy under the Cloud objects tab.
- 2 Enter a value for this parameter in MB: `COS_SHM_BUFFER_SIZE =`
For example: `COS_SHM_BUFFER_SIZE = 16`

Configuring the number of buffers

By default, NetBackup creates seven buffers for each read or write operation. Depending on the availability of memory on the backup host, you can set the number

of buffers to 4 to 16 buffers for each stream. Increasing the number of buffers results in faster backups, but the memory usage on the backup host may increase.

To configure the number of buffers:

- 1 Open the `/usr/opensv/netbackup/bp.conf` file on the backup host.
- 2 Enter a value for this parameter in MB: `COS_NO_SHM_BUFFER =`
For example: `COS_NO_SHM_BUFFER = 12`

Note that these settings are applied to all backup jobs that use that backup host.

Configure a temporary staging location

This configuration is applicable only for Cloud object store protection using dynamic multi-streaming and backup hosts of NetBackup version 11.0 or later.

Enhanced Dynamic multi-streaming backup in NetBackup 11.0, or later, uses a staging location to exchange object data between the object store reader and NetBackup data mover. This staging location must be a filesystem path on the Cloud object store backup host. The storage space on the staging location is reclaimed during the course of a backup as and when the object data and metadata are moved to the backup target. You can configure the storage space and water mark value depending upon the object store environment that you are protecting.

For more information about configuring the temporary storage location, see the Knowledge article:

<https://support.cohesity.com/s/article/article-100073852.html.html>

Default paths:

- For BYO default path is: `/usr/opensv/netbackup/db/cos_tmp_staging_path`.
- For Cloud scale primary and media server: The default temporary staging location: `/usr/opensv/netbackup/db/cos_tmp_staging_path` is symlinked to mounted location
`/mnt/nbdata/usr/opensv/netbackup/db/cos_tmp_staging_path`
- For Flex:
 - The default temporary staging location for a media server:
`/usr/opensv/netbackup/db/cos_tmp_staging_path` is symlinked to the mounted location
`/mnt/nbstage/usr/opensv/netbackup/db/cos_tmp_staging_path`
 - For the primary server:
The temporary staging path
`/usr/opensv/netbackup/db/cos_tmp_staging_path` is symlinked to the

mounted location

```
/mnt/nbdata/usr/opensv/netbackup/db/cos_tmp_staging_path
```

For backup host other than Flex, it is recommended to use different path than the default path as it may not comply with the storage requirements. It is also recommended to create a different mount point with enough space. Backup performance is affected if the minimum required space is not available.

Note: Ensure that the *automount* configuration entry is updated on the backup host, if you are using a different path other than the default.

Configuring the temporary staging location:

- Use the following guidelines for sizing of the temporary staging location on a backup host:
 - The minimum space requirement is 20 GB per backup job per bucket or container being backed up by the backup host.
 - For increased backup speed, you can configure 100 GB to 500 GB space for staging.
- The device used for the temporary staging location must have high read and write throughput. It is recommended to use an SSD. If you still experience slow disk I/O throughput, mount RAM as a disk, and configure as the temporary staging location.

Here is an example command to mount RAM as disk:

```
sudo mount -t tmpfs -o size=64G tmpfs /mnt/tmp
```

For 64 GB and 128-GB RAM setups, it is recommend mounting half of RAM as a disk. For a 32-GB RAM setup, it is not recommended to mount RAM as a disk, as low backup performance is expected on large size objects.

- Ensure that there is enough free space at the location.
- File path should have only ASCII Characters.
- Ensure that the NetBackup service user is the owner of the storage location and has 0700 permissions. Writing fails if the permissions are set to anything other than 0700.
- The temporary staging location is cleaned up once a backup completes.
- Each backup host must have a different mount point for the temporary staging location.

Optionally, you can set the following parameters in the `bp.conf` file in each backup host, to configure the temporary staging location.

Configuring advanced dynamic multi-streaming parameters

You can configure these optional parameters in the backup host.

Table 2-2 Advanced parameters for backup host

Parameter	Default/ Configuration	Description
COSP_LIST_API_MAX_OBJECTS	Default: 5000 for Azure backups and 1000 for AWS/S3 backups. Configuration: 100 - 5000 for Azure 100 - 1000 for S3/AWS	During object listing, this value is used as the batch size in the ListObject API. Reducing this value may slow down the overall backup process. It is recommended to set it according to your cloud provider's specifications or recommendations.
COSP_DOWNLOAD_OBJECT_PART_SIZE_IN_MB	Default: 16 Configuration: 1 - 512	Represents the size of each chunk or part used when downloading large objects in multiple chunks. If the cloud is fast and capable of handling larger chunks efficiently, increasing this value may improve the overall backup speed.
COSP_TOTAL_OBJECT_DOWNLOAD_WORKERS	Default: 160 Configuration: 10 - 1000	This parameter determines the number of parallel tasks used to download objects from a specified bucket. If the cloud is slow, increasing this value can cause issues such as request failures or retries. However, if the cloud is fast, increasing this value can speed up the backup process. It's also important to note that increasing the number of workers consumes more CPU resources. Since this value applies to each backup job, it's crucial to configure it carefully, especially when running multiple parallel jobs.

Table 2-2 Advanced parameters for backup host (*continued*)

Parameter	Default/ Configuration	Description
COSP_STREAM_CHANNEL_BUFFER_SIZE	Default: 100 Configuration: 100 - 5000	This parameter restricts the total buffer size/number of objects to keep in the cache for processing after listing. If multiple streams are running, then this is the buffer/cache size of each stream.
COSP_DB_INSERT_BATCH_SIZE	Default: 500 Configuration: 500 - 5000	Represents the number of downloaded objects that are sent to bpbkar for final processing. Increasing this value to an optimistic number, such as 1000, can boost the overall backup speed, provided the staging location has sufficient space to handle the batch size. However, if the staging location lacks enough space, increasing this value may slow down the backup process.
COSP_STAGING_LOC_WATER_MARK_IN_PERCENTAGE	Default: 80 Configuration: 30 - 80	This parameter restricts the usages of the staging location. It accepts the value in percentage. The default value is 80. The range is 30 to 80. For example, COSP_STAGING_LOC_WATER_MARK_IN_PERCENTAGE= 70
COSP_SPACE_MANAGEMENT_TIMEOUT_IN_MIN	Default: 5 Configuration: 5 - 60	The time to wait before starting a backup, if the required space is not available on the temporary staging location. The value is in minutes. Default value = 10. The range is 5 to 60 (minutes). For example: COSP_SPACE_MANAGEMENT_TIMEOUT_IN_MIN = 20

You can configure this optional parameter in the primary server.

Table 2-3

Parameter	Default/ Configuration	Description
DYNAMIC_STREAMING _START_CHILD_BACKUP _JOBS_TIMEOUT	Default: 600 (Seconds) Configuration: 600 - 3600 (Seconds)	This timeout value determines how long each parent job waits to start all its child jobs, when multiple jobs are triggered with multiple streams. If the timeout is reached, the job fails with a crawler timeout error. We recommend setting this value as high as possible if multiple backups are running. Note that this setting takes effect only before running a backup job.

Permissions required for Amazon S3 cloud provider user

The 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
- s3:PutObjectRetention
- s3:BypassGovernanceRetention
- s3:GetBucketObjectLockConfiguration
- s3:Getobjectretention

To use the object change tracking feature, add the following additional permission for the source buckets:

- s3:GetBucketLogging
- s3:GetBucketAcl

To use the object change tracking feature, add only the following permissions in the log buckets:

- s3:ListBucket
- s3:GetObject
- s3:GetBucketAcl

Permissions required for Azure blob storage

Here is a custom role definition (in JSON format) that is required for the discovery, backup, restore, and authentication of the Microsoft Azure object store. You must associate a custom role with these permissions that the NetBackup users can use to work with Azure blob. To use Service principal or Managed identity authentication you need a role with the following permissions.

```
{
  "properties": {
    "roleName": "cosp_minimal",
    "description": "minimal permission required for cos
protection.",
    "assignableScopes": [
      "/subscriptions/<Subsfriction_ID>"
    ],
    "permissions": [
      {
        "actions": [

          "Microsoft.Storage/storageAccounts/blobServices/read",

          "Microsoft.Storage/storageAccounts/blobServices/containers/read",

          "Microsoft.Storage/storageAccounts/blobServices/containers/write",
          "Microsoft.ApiManagement/service/*",
          "Microsoft.Authorization/*/read",

          "Microsoft.Resources/subscriptions/resourceGroups/read",
          "Microsoft.Storage/storageAccounts/read"
```

```

    ],
    "notActions": [],
    "dataActions": [

        "Microsoft.Storage/storageAccounts/blobServices/containers/blobs/filter/action",

        "Microsoft.Storage/storageAccounts/blobServices/containers/blobs/tags/read",

        "Microsoft.Storage/storageAccounts/blobServices/containers/blobs/tags/write",

        "Microsoft.Storage/storageAccounts/blobServices/containers/blobs/write",

        "Microsoft.Storage/storageAccounts/blobServices/containers/blobs/read",

        "Microsoft.Storage/storageAccounts/blobServices/containers/blobs/immutableStorage/immutableStorage/action",

    ],
    "notDataActions": []
  }
}
}
}
}

```

Permission required for Azure Data Lake Storage

NetBackup protection workflows—such as asset discovery, backup, and restore—for Azure Data Lake Storage (ADLS) can be authenticated using Azure user managed or system managed identities. To use a managed identity for ADLS, you must:

- Assign the 'Storage Blob Data Owner' role to the managed identity to enable ACL restoration.
- Assign a role to the managed identity that includes the following permissions:

```

{
    "id": "/subscriptions/<subscription
id>/providers/Microsoft.Authorization/roleDefinitions/<role id>",

```

```

"properties": {

    "roleName": "cosp_minimal",

    "description": "minimal permission required for cos
protection.",

    "assignableScopes": [

        "/subscriptions/<subscription id>"

    ],

    "permissions": [

        {

            "actions": [

                "Microsoft.Storage/storageAccounts/blobServices/read",

                "Microsoft.Storage/storageAccounts/blobServices/containers/read",

                "Microsoft.Storage/storageAccounts/blobServices/containers/write",

                "Microsoft.ApiManagement/service/*",

                "Microsoft.Authorization/*/read",

                "Microsoft.Resources/subscriptions/resourceGroups/read",

                "Microsoft.Storage/storageAccounts/read",

                "Microsoft.Storage/storageAccounts/blobServices/containers/setAcl/action"
            ]
        }
    ]
}

```

```

    ],

    "notActions": [],

    "dataActions": [

        "Microsoft.Storage/storageAccounts/blobServices/containers/blobs/filter/action",

        "Microsoft.Storage/storageAccounts/blobServices/containers/blobs/tags/read",

        "Microsoft.Storage/storageAccounts/blobServices/containers/blobs/tags/write",

        "Microsoft.Storage/storageAccounts/blobServices/containers/blobs/write",

        "Microsoft.Storage/storageAccounts/blobServices/containers/blobs/read",

        "Microsoft.Storage/storageAccounts/blobServices/containers/blobs/add/action",

        "Microsoft.Storage/storageAccounts/blobServices/containers/blobs/immutableStorage/runAsSuperUser/action"

    ],

    "notDataActions": []

}

]

```

```
}  
  
}
```

Permissions required for GCP

GCP requires the following permissions to work with NetBackup:

```
storage.bucketOperations.cancel  
storage.bucketOperations.get  
storage.bucketOperations.list  
storage.buckets.create  
storage.buckets.createTagBinding  
storage.buckets.delete  
storage.buckets.deleteTagBinding  
storage.buckets.enableObjectRetention  
storage.buckets.get  
storage.buckets.getIamPolicy  
storage.buckets.getObjectInsights  
storage.buckets.list  
storage.buckets.listEffectiveTags  
storage.buckets.listTagBindings  
storage.buckets.restore  
storage.buckets.setIamPolicy  
storage.buckets.update  
storage.multipartUploads.abort  
storage.multipartUploads.create  
storage.multipartUploads.list  
storage.multipartUploads.listParts  
storage.objects.create  
storage.objects.delete  
storage.objects.get  
storage.objects.getIamPolicy  
storage.objects.list  
storage.objects.restore  
storage.objects.setIamPolicy  
storage.objects.update
```

Prerequisites for GCP protection

NetBackup supports GCP backup using Amazon S3-standard API. To protect GCP buckets you must enable S3 interoperability settings on GCS or Google Cloud Project.

Limitations and considerations

Consider the following when protecting Cloud object store workloads.

- NetBackup does not allow a prefix or object query that starts with "/". Examples:

```
prefix = /  
prefix = /folder1  
prefix = /object1  
prefix = folder1//  
object = /obj1
```

- NetBackup does not back up an object with its name in the format: <name>/
- NetBackup does not support backups of blobs or objects that are 0 bytes in size and contain a slash ("/") in their names. For example, `TestObject/`.
- If the name field of an object, blob, or folder contains only slashes (/) and no other characters, the name appears empty in the folder hierarchy in the NetBackup UI's restore view.
- Checkpoint restart is not supported for policies having Object change tracking enabled.
- Checkpoint restart with dynamic multi-streaming is supported in NetBackup 11.2 and later. Ensure that the backup host, media server, deduplication backup host, storage unit, and storage unit group are all running NetBackup version 11.2 or later.
- You can only protect entire buckets with Dynamic multi-streaming.
- You must configure backup hosts with at least the minimum recommended staging space.
- For Cloud Object Store backups in a NetBackup Flex Scale environment, the staging location on the media server pod provides approximately 45 GB of available space. If the total size of the cloud objects to be backed up exceeds the available free space in the staging area, the backup job may fail.

Limitations based on the upgrade scenario

- Upgrading to NetBackup version 11.2:
 - It is recommended to update the primary server, media server, and backup hosts to the current version.
 - Upgrade only the primary server: Backup jobs with Dynamic multi-streaming fails. You must update the backup host and the media server used in the policy to the current version.

- Upgrade only the primary and media servers: Backup jobs with Dynamic multi-streaming fails. You must update the backup host used in the policy to the current version.
- Upgrade only primary server and backup host and media server remains at an older version: NetBackup protection works.
- To enable Checkpoint restart in a Dynamic multi-streaming enabled policy, upgrade the primary, media, and the backup host to NetBackup version 11.2 or later.
- If you upgrade to the latest NetBackup version from version 10.1 or 10.2, the following limitations occur:
 - You can create Cloud object store accounts with backup hosts or scale-out servers of version 10.3 or later only. You cannot update an existing Cloud object store account that was created on NetBackup 10.3 or later with backup hosts or scale-out servers older than version 10.3.
 - You can create policies only with backup hosts or scale-out servers of version 10.3 or later. You cannot update an existing policy that is created on NetBackup 10.3 or later with backup hosts or scale-out servers older than version 10.3.
 - The following credential types are not supported with backup hosts or scale-out servers older than version 10.3: For Azure: Service principal and Managed identity. For AWS: Assume role (EC2).
 - Restores with object lock properties are supported for backup hosts or scale-out servers of version 10.3 or later only.
 - Backup and restore of buckets with default retention enabled are supported with backup hosts or scale-out servers of version 10.3 or later only.
- For Azure, if you update a policy created with NetBackup version prior to 10.3, with a backup host or scale-out server of version 10.3 or later, the backups fail. As a workaround, update all the buckets to use the new format of the provided generated ID with the existing queries. Note that you must create the associated Cloud object store account in the policy, using NetBackup 10.3 or later, for this workaround to be successful.
- Discovery is supported for NetBackup version 10.3 or later, deployed on RHEL, and NetBackup version 11.0.0.1 or later deployed on RHEL and SUSE. If no supported host is available, then discovery does not start for any of the configured Cloud storage accounts. In this case, discovery status is not available, and you cannot see a bucket list during policy creation. Even if you add the buckets manually after discovery fails, your backups may fail. Upgrade at least one supported backup host or scale-out server and create a new policy.

- If you update a policy that is created on a NetBackup version prior to 10.3, consider the following after a backup:
 - After backup, you may see two versions of the same buckets, for the old and new formats. If you want to restore old data, select the bucket in the old format. For newer backups, select the ones in the newer format.
 - The subsequent backup after the update is a full backup, irrespective of what is configured in the policy.
- When you upgrade to 10.3, the first Azure blob accelerated backup takes a backup of all objects in the selection, even if the configured backup is incremental. This full backup is required for the change in metadata properties for the Azure blobs between NetBackup versions 10.2 and 10.3. The subsequent incremental backups back up only the changed objects.
- If you use a Cloud object store account created in a version older than 10.3, NetBackup discovers the buckets with the old format, where:
`uniqueName=bucketName`.

Adding Cloud object store accounts

Adding a Cloud object store account is the first step in protecting a workload. You can add one or more accounts in a NetBackup primary server. 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, for every 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 credential to be used for the set of buckets or containers, you can create multiple accounts.

The required permissions for backup and recovery are different. See if it is helpful to create separate accounts for backup and recovery. Use something other than the original bucket options to restore to a different Cloud object store account during recovery.

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

For Cloud object store accounts, 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.

You need to select a validation host while creating a Cloud object store account. A validation host is a specific backup host that validates the credentials. The validation host is used during manual, periodic discovery, and when manual validation is required for an existing Cloud object store account. The validation host can be different from the actual backup host specified in the policy.

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**. Enter a name for the account in the **Cloud object store name** field, and select a provider from the list **Select Cloud object store provider**.
- 3** To select a backup host or scale-out server, click **Select host for validation**. The host must be NetBackup version 10.1 or later, on a RHEL media server, or NetBackup 11.0.0.1 or later, on a RHEL or SUSE media server. Only these servers support credential validation, backup, and recovery in the Cloud object stores.
 - To select a backup host, select the **Backup host** option, and select a host from the list.
 - To use a scale-out server, select the **Scale out server** option, select a server from the list. NetBackup Snapshot Manager servers 10.3 or later, serve as scale-out servers.

If you have a very large number of buckets, you can also use NetBackup Snapshot Manager as a backup host with NetBackup 10.3 or later releases. Select the **Scale out server** option, and select a NetBackup Snapshot Manager from the list.

Note: Your existing NetBackup primary server must be already configured with this instance of NetBackup Snapshot Manager.

Note: In a Cloud Scale environment, you cannot use the primary server as a backup host. For more information regarding the Cloud Scale environment, see *NetBackup Kubernetes Administrator's guide*.

- 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 38.. 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 the region during account creation to list the bucket.

- 5 In the **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 or scale-out server must be hosted on the EC2 instance. Make sure the IAM role associated with the EC2 instance has required permissions to access the required cloud resources for Cloud object store protection. Make sure that you select the correct region as per permissions associated with the 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**-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 to 10 as required to create the account.
See [“Creating cross-account access in AWS ”](#) on page 34.
- **Assume role (EC2)**- NetBackup retrieves the AWS IAM role credentials that are associated with the selected backup host or scale-out server, hosted on an EC2 instance. Henceforward, NetBackup assumes the role mentioned in the Role ARN to access the cloud resources required for Cloud object store protection.
- **Credentials broker**- NetBackup retrieves the credentials to access the required cloud resources for Cloud object store protection.
- **Service principal**- NetBackup uses the tenant ID, client ID, and client secret associated with the service principal to access the cloud resources required for Cloud object store protection. Supported by Azure.
- **Managed identity**- NetBackup retrieves the Azure AD tokens, using the managed identity that is associated with the selected backup host or

scale-out server or the user. NetBackup uses these Azure AD tokens to access the required cloud resources for Cloud object store protection. You can use system or user-assigned managed identities.

- 6 You can add existing credentials or create new credentials for the account:
 - To select an existing credential for the account, select the **Select existing credentials** option, select the required credential from the table, and click **Next**.
 - To use **Managed identity** for Azure, select **System assigned** or **User assigned**. For the user-assigned method, enter the **Client ID** associated with the user to access the cloud resources.
 - 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:
 - For the **Access key** method, provide **Storage account** credentials, specify **Storage account**.
 - For the **Service principal** method, provide **Client ID**, **Tenant ID**, and **Secret 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 the 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 cloud providers, NetBackup provides the capability to verify the SSL certificates revocation status by OCSP protocol. The OCSP protocol sends a validation request to certificate issuer to get certificates current revocation status. If SSL is enabled and the check certificate revocation option is enabled, each non-self-signed SSL certificate is verified with an OCSP

request. 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 SSL mode. Ensure that the cloud server (public or private) has a CA-signed certificate. If it does not have the CA-signed certificate, data transfer between NetBackup and the cloud provider fails in SSL mode. If you want to work with self-signed SSL certificate, the certificate has to be added in NetBackup's cloud storage CA trust store. See [“Managing Certification Authorities \(CA\) for NetBackup Cloud ”](#) on page 35.

Note: The FIPS region of the 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 a 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 the HTTP proxy type.

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

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

- **None**—Authentication is not enabled. A username and password are not required.
- **Basic**—Username and password needed.
- **NTLM**—Username and password needed.

Username—is the username of the proxy server.

Password—can be empty. You can use a maximum of 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 or scale-out server 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 is deployed in one of these accounts, you can protect your data in all your AWS accounts. You need to configure cross-account data access in the AWS portal, before you select **Assume role** or **Assume role EC2** 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 that you want to protect.
- 3** Create a new policy for the IAM role and ensure that it has the required permissions to access the bucket and objects in that target AWS account. See [“Permissions required for Amazon S3 cloud provider user”](#) on page 21.
- 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 the capability to verify the revocation status of SSL certificates using the Online Certificate Status Protocol (OCSP). If SSL and the **Check certificate revocation** option are both enabled, NetBackup verifies each SSL certificate. To verify, NetBackup makes an OCSP request to the CA to check the revocation status of the certificate presented during the SSL handshake. NetBackup does not connect to the cloud provider, if the status is returned as revoked or it fails to connect to the OCSP endpoint present in the SSL certificate.

To enable validation, update the `Check certificate revocation` property from the Cloud object store account dialog.

Requirements for enabling the Check certificate revocation option

- OCSP endpoints are HTTP thus, turn off any firewall rule that blocks HTTP (port 80) connections to external networks. 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, the OCSP URL's endpoint supports 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.
- The OSCP URL of the CA should be present in the certificate's Authority Information Access extension.

Managing Certification Authorities (CA) for NetBackup Cloud

NetBackup 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 the following location:

- **Windows:**

```
<installation-path>\NetBackup\var\global\cloud
```

- **UNIX:**

```
/usr/opensv/var/global/cloud/
```

Note: In a cluster deployment, the 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 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

- Starfield Services Root Certificate Authority - G2
- Baltimore CyberTrust Root
- DigiCert Assured ID Root CA
- DigiCert Assured ID Root G2
- DigiCert Assured ID Root G3
- DigiCert Global CA G2
- DigiCert Global Root CA
- DigiCert Global Root G2
- DigiCert Global Root G3
- DigiCert High Assurance EV Root CA
- DigiCert Trusted Root G4
- D-Trust Root Class 3 CA 2 2009
- GlobalSign Root CA
- GlobalSign Root CA - R3
- COMODO RSA Certification Authority
- AAA Certificate Services
- Go Daddy Root Certificate Authority - G2
- ISRG Root X1

Adding a new region

You can add new regions for a 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 selection is not available for some Cloud object store providers.

The option to add a new region is not available in the case of Azure Data Lake Storage and Azure Data Lake Storage Government provider types.

To add a region:

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

For the cloud providers that support the 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 the location constraint as blank, use 'us-east-1' as the location constraint.

- 2 Enter the service URL. For example: `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 on 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.

Manually discovering assets for protection

To discover assets manually

- 1 On the left, click **Workloads > Cloud object store**.
- 2 Select the row of the account for which you want to run discovery, and click **Discover** at the top. Or, click the ellipsis menu (three dots) in the row of the account, and click **Discover**.

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 29.

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**.

Scan for malware

NetBackup version 10.5 and later provides support for scanning Cloud object store assets for malware through the Cloud-Object-Store policy type.

For triggering malware scan, 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*.

Backup images

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

To scan policy of client backup images for malware

- 1 On left, click **Detection and reporting > Malware detection**.
- 2 On the **Malware detection** page, click **Scan for malware**.
- 3 In the **Search by** option, select **Backup images**.
- 4 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**: Select the policy type as **Cloud-Object-Store**.
 - **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 (DataDomain) and AdvancedDisk storage type disk pools are listed.
 - **Disk type**: MSDP (PureDisk), OST (DataDomain) and AdvancedDisk disk types are listed.
 - **Malware scan status**

- For the **Select the timeframe of backups**, verify the date and the time range or update it.
- 5 Click **Search**: Select the search criteria and ensure that the selected scan host is active and available.
- 6 From the **Select the backups to scan** table select one or more images for scan.
- 7 In the **Select a malware scanner host pool**, **Select** the appropriate host pool name.

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

- 8 Click **Scan for malware**.
- 9 After the scan is initiated, the **Malware Scan Progress** is displayed.

The following are the status fields:

- **Not scanned**
- **Not infected**
- **Infected**
- **Failed**

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

Note: Any backup images that fail validation are ignored. Malware scanning is supported for the backup images that are stored on storage with instant access capability and for the supported policy types only.

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

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

Assets by policy type

This procedure is not applicable for YARA scanning.

The following section describes the procedure for scanning Cloud object store assets for malware through the Cloud-Object-Store policy type.

To scan the supported assets by policy type, perform the following:

- 1** On the left, click **Detection and reporting > Malware detection**.
- 2** On the **Malware detection** page, click **Scan for malware**.
- 3** In the **Search by** option, select **Assets by policy type**.
- 4** From the **Client/Asset** table, select a Client/Asset to scan.
- 5** Click **Next**.
If the selected client in the previous step supports multiple policy types, you can select a single policy type for scanning.
- 6** For the **Start date/time** and **End date/time** verify the date and the time range or update it.
The scan is initiated for a maximum of 100 images.
- 7** In the **Scanner host pool**, select the appropriate host pool name.
- 8** 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**
- 9** Click **Scan for malware**.

Warning: Scan is limited to only 100 images. Adjust the date range and try again.

- 10** After the scan is initiated, the **Scan status** is displayed. The following are the status fields:
 - **Not scanned**
 - **Not infected**
 - **Infected**
 - **Failed**

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

Any backup images that fail validation are ignored. Malware scanning is supported for the backup images that are stored on storage with instant access capability and for the supported policy types only.

- **Pending**
- **In progress**
- **Aborted on detection of malware**

For more information on the malware scan status, refer to the *NetBackup Security and Encryption Guide*.

Protecting Cloud object store assets

This chapter includes the following topics:

- [About accelerator support](#)
- [About incremental backup](#)
- [About dynamic multi-streaming](#)
- [Enable auto-resume for incomplete backup jobs](#)
- [About object change tracking](#)
- [About storage lifecycle policies](#)
- [About policies for Cloud object store assets](#)
- [Planning for policies](#)
- [Prerequisites for Cloud object store policies](#)
- [Creating a backup policy](#)
- [Policy attributes](#)
- [Creating schedule attributes for policies](#)
- [Configuring the Start window](#)
- [Configuring the exclude dates](#)
- [Configuring the include dates](#)
- [Configuring the Cloud objects tab](#)

- [Adding conditions](#)
- [Adding tag conditions](#)
- [Examples of conditions and tag conditions](#)
- [Managing Cloud object store policies](#)

About accelerator support

NetBackup accelerator for Cloud object store optimizes the backups. The backup host or scale-out server uses the change detection techniques to determine the current state of the Cloud object store's objects or blobs to identify the changes that occurred since the last backup. The backup host or scale-out server 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. For example, the backup host, during Cloud object store policy execution.
- Creates a compact backup stream that uses less network bandwidth between the backup host or scale-out server and media server.
- Creates a backup image that contains all data that is needed for restoration.

How NetBackup accelerator works with Cloud object store

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

- If the backup host or scale-out server 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 or scale-out server 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: <https://support.cohesity.com/s/article/article-100040093>
- Storage unit groups are supported only if the storage unit selection in the group is Failover.
- Supports 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 47.
- If a previous backup of the policy, bucket and query does not exist on the backup host or scale-out server, NetBackup performs a full backup, and creates a track log on the backup host or scale-out server. This initial backup occurs at the speed of a normal (not accelerated) full backup. Subsequent Accelerator backups using the same backup host or scale-out server 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 be completed 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 to 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 have 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 changes in metadata, NetBackup uses one or more cloud APIs per object/blob. Hence, change detection time increases with the number of objects/blobs to be processed. You may observe backups running longer than expected in cases with little or no data change but a large number of objects.
- If in your environment, for a given object, the metadata or tag is always changed (added/removed/updated) with its data. Evaluate using incremental without accelerator over incremental with accelerator from a performance and cost viewpoint.
- While creating a Cloud object store policy with multiple tag-based queries, you can use a few simple rules to get the best effect with accelerator. Use the query builder in the policy creation page, and create separate queries, one query per tag. The accelerator-based policies perform best in this configuration. NetBackup has two ways to determine whether an object is considered for incremental backup or not.
 - The object's modification time.
 - Any changes in the tags and user attributes.

These metadata checks are not required for the organization environments where the metadata of the objects do not change after the objects are created. You can use the **Quick Object change scan** option in the policy to avoid these metadata checks, which lead to faster change detection and faster accelerated backup.

Accelerator force rescan for Cloud object store (schedule attribute)

Accelerator force rescan is a property of the full backup schedule.

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 fingerprints the data and uses 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 or scale-out server and the server.

Accelerator backup and NetBackup catalog

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; 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 one, replacing incremental backups with full backups increases the catalog size.

Calculate the NetBackup accelerator track log size

The accelerator track log stores object metadata and file fingerprints in 128-KB segments. The track log is stored on the backup host. The track log size is relative to the size and number of objects being backed up. Track logs are created separately for each policy, backup selection (cloud account and bucket), and stream combination.

The following are only guidelines; requirements in a specific environment may differ. In environments where data changes often, a bigger track log size may be required.

You can use the following formula to calculate the approximate track log size:

```
Track log Size in Bytes = 2 * ( (Number of objects * 200) + ((Total  
used disk space in KiB/128KiB) * 20) )
```

For example: A 1-TB file system with one million objects needs approximately 701-MB track log.

Note that if you modify the backup selection or stream count in an accelerator-enabled policy, NetBackup creates a new track log. The older track logs remain on the backup host.

About incremental backup

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

For Cloud object store workload, some metadata properties 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 as a loss of data during an incremental backup.

For Azure Data Lake and Azure Data Lake Government providers, when you update the ACLs of files or directories, the last modified time of the file or directory doesn't change. So, only by changing the ACLs, the files and directories do not qualify as incremental backups.

For a 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 a 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 the last backup (either full or last incremental, based on the 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 dynamic multi-streaming

Multi-streaming backup for Cloud object store policy runs simultaneous backup streams for a given backup selection. The backup selection is divided into several streams that run in parallel, resulting in a faster backup. The number of streams can be configured for each policy in the backup selection tab of the cloud object store policy. Each backup stream creates a unique backup image. Eventually, all the images created by the streams for that backup selection represent the backup of that specific selection.

Dynamic multi-streaming is enabled by default on all newly created Cloud object store policies.

Specifying the maximum number of streams

You can specify the maximum number of streams that you want to use for a bucket or container in the policy attributes.

See [“Policy attributes”](#) on page 68.

Considerations for using dynamic multi-streaming

- Entire buckets/containers are backed up when you use dynamic multi-streaming.
- The number of streams that you specify in a policy is applicable for each of the buckets that the policy protects. For example, if you specify 10 streams in the policy and select five buckets for backup, you get 50 concurrent streams. Some streams may go to a queue, if the maximum number of concurrent jobs allowed in the storage unit selected for the policy, is less than the total number of streams that are running across different policies. For optimal performance, keep the **Maximum concurrent jobs** allowed property of selected storage greater than the total number of streams that you expect to run across the policies.
- You cannot use a scale-out server as a backup host, when you use dynamic multi-streaming.
- Dynamic multi-streaming requires staging location path.
- Job retry feature does not work for backup jobs.
- Dynamic multi-streaming is not available for Azure Data Lake storage and Azure Data Lake storage Government providers.
- You must configure a temporary staging location as per guidelines to use this feature. See [“Configure a temporary staging location”](#) on page 17.
- Dynamic multi-streaming starts all the backup streams for a bucket or container at the same time and writes to a storage unit. Therefore, using tape storage units as the target for primary backup copies is not recommended. You can use an MSDP storage as the target for the first backup copy, and configure a tape storage as the target for secondary or duplication copies.

Enable auto-resume for incomplete backup jobs

With this feature, when a backup job that has checkpoint restart enabled enters an incomplete state, the job automatically resumes after a configured time interval. If the resumed job fails again, the NetBackup marks it as incomplete until a scheduled retry runs. NetBackup marks the job as failed after all scheduled retry attempts have completed. You can configure the number of retries for each job and the delay between retry attempts.

For details of the Checkpoint restart attribute: See [“Policy attributes”](#) on page 68.

To configure auto-resume:

- 1 On the left, click **Host properties**, under **Hosts**.
- 2 Select the host that you want to configure. If the host is not connected, click the ellipsis menu (three dots) in the row of the host, and click **Connect**. After the host connects, click the ellipsis menu (three dots) in the row of the host, and click **Edit primary server**.
- 3 Click **Global attributes**.
- 4 To set the interval between retries, specify a value in minutes in the **Job retry delay** field.
- 5 In the **Schedule backup attempts** fields, specify the following:
 - How many preceding hours NetBackup examines to determine whether to allow another backup attempt (retry). By default, NetBackup examines the past 12 hours.
 - How many times a backup can be retried within that timeframe. By default, NetBackup allows two attempts. Attempts include the scheduled backups that start automatically or the scheduled backups that are user-initiated.

For example, consider the following example scenario using the default setting 2 tries every 12 hours:

- Policy_A runs at 6:00 P.M.; Schedule_1 fails.
- Policy_A is user-initiated at 8:00 P.M.; Schedule_2 fails.
- At 11:00 P.M., NetBackup looks at the previous 12 hours. NetBackup sees one attempt at 6:00 P.M., and one attempt at 8:00 P.M. The **Schedule backup attempts** setting of two has been met so NetBackup does not try again.
- At 6:30 A.M. the next morning, NetBackup looks at the previous 12 hours. NetBackup sees only one attempt at 8:00 P.M. The **Schedule backup attempts** setting of two has not been met so NetBackup tries again. If a schedule window is not open at this time, NetBackup waits until a window is open.

About object change tracking

NetBackup 11.1 introduces Object change tracking, a feature that leverages object storage vendors' native mechanisms to identify the changed objects. S3 Bucket Logging is one such feature that tracks operations such as create, update, and delete. NetBackup analyzes these logs to efficiently identify changes and perform efficient incremental backups.

Advantages of object change tracking for incremental backup

Compared to traditional object scanning, object change tracking offers several advantages in detecting changed objects and performing incremental backups more efficiently. This option is ideal for buckets with many objects and low change rates.

- NetBackup does not need to scan every object in the bucket to check for changes since the last backup. This enables faster backups.
- Reduced API costs.
- By eliminating full object scans, storage endpoints avoid heavy API traffic, leading to improved overall performance.
- Incremental backups become scalable as change detection is independent of the number of objects in the bucket or the size of the bucket.

Configuring object change tracking

Complete the configuration steps in the subsequent sections before using object change tracking for IBM Storage Ceph and Amazon S3.

Since NetBackup object change tracking uses the bucket logging feature in IBM Storage Ceph, you must enable and configure bucket logging in IBM Storage Ceph.

IBM Storage Ceph tracks change events on the source bucket and logs them in a log bucket. Log buckets function like any other bucket, but you must grant write access from different source buckets.

Configuring access permissions for the buckets

"You must configure access permissions for the Amazon S3 provider to enable communication for the log and source buckets.

See ["Permissions required for Amazon S3 cloud provider user"](#) on page 21.

Configuring access policy on the log bucket

You must configure a policy on the log bucket to enable logging for all other buckets. Here is a sample policy document for the log bucket. The policy name is *log-bucket*, and *source-bucket* as the source bucket name.

```
{  
  
  "Version": "2012-10-17",  
  
  "Statement": [  

```

```
{  
  
  "Sid": "AllowLoggingFromSourceBucket",  
  
  "Effect": "Allow",  
  
  "Principal": {  
  
    "Service": "logging.s3.amazonaws.com"  
  
  },  
  
  "Action": "s3:PutObject",  
  
  "Resource": "arn:aws:s3:::log-bucket/prefix*",  
  
  "Condition": {  
  
    "ArnLike": {  
  
      "aws:SourceArn": "arn:aws:s3:::source-bucket"  
  
    }  
  
  }  
  
}
```

Configuration guidelines for IBM Storage Ceph

Refer to IBM Storage Ceph documentation for details of any of these points.

- You must enable bucket logging on the source buckets with the required configuration. See [“Enable bucket logging for source buckets”](#) on page 54.
- You must create the log bucket before enabling bucket logging in the Ceph console.
- Create a policy for the log bucket so that the log bucket can be accessible to the source bucket logging. You must create the log bucket policy before enabling

bucket logging, otherwise the operations on source buckets fail. See [“Configuring access policy on the log bucket”](#) on page 52.

- Both the source and log buckets must be accessible using the same credentials defined in the Cloud object store account in NetBackup.
- Do not use the log bucket as source bucket in bucket logging.
- All source buckets must have a unique prefix. The prefix cannot be empty.
- Log object key format must be Date partition. The `PartitionDateSource` must be configured as `EventTime`.
- To delete the older log objects, create a bucket policy. See [“Maintaining the log bucket”](#) on page 59.

Enable bucket logging for source buckets

Refer to the IBM Storage Ceph documentation to enable bucket logging on the source bucket using the Ceph console. This section explains how to enable bucket logging using the S3 API.

Method: PUT

API Endpoint:

`http://<IP-of-CEPH-server>:<CEPH-RGW-Port>/<source-bucket-name>?logging`

Where:

- *IP-of-CEPH-server*—IP address of the Ceph server.
- *CEPH-RGW-Port*—Port of the S3 gateway.
- *source-bucket-name*—the source bucket on which you want to enable bucket logging.

API payload:

```
<?xml version="1.0" encoding="UTF-8"?>
```

```
<BucketLoggingStatus xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
```

```
  <LoggingEnabled>
```

```
    <TargetBucket><target-bucket-name></TargetBucket>
```

```
    <TargetPrefix><target-prefix></TargetPrefix>
```

```
    <ObjectRollTime>300</ObjectRollTime>
```

```
<LoggingType>Journal</LoggingType>

<RecordsBatchSize>0</RecordsBatchSize>

<TargetObjectKeyFormat>

    <PartitionedPrefix>

        <PartitionDateSource>EventTime</PartitionDateSource>

    </PartitionedPrefix>

</TargetObjectKeyFormat>

</LoggingEnabled>

</BucketLoggingStatus>
```

Where:

- *<target-bucket-name>* is the bucket where target logs are generated.
- *<target-prefix>* defines the prefix (path) that Amazon S3 uses to store log files in the target bucket. For example, the logs may be stored as:
`s3://target-bucket-name/targetPrefix/2025-10-28-12-34-56-123456789ABCDEF`
- *ObjectRollTime* specifies how often the logging system rotates log files—closing the current log object and starting a new one.
- *LoggingType* specifies the format or mode used to write logs to the target bucket. Must be *Journal*.
- *RecordsBatchSize* specifies the number of log records (events) the system collects before writing them to the target log object.
- *PartitionDateSource* specifies which timestamp the logging system uses to create date-based partitions (subfolders) for log objects in the target bucket. Must be *EventTime*.

Verify bucket logging using S3 API

Method: GET

API Endpoint:

`http://<IP-of-CEPH-server>:<CEPH-RGW-Port>/<source-bucket-name>?logging`

Where:

- *IP-of-CEPH-server*—IP address of the Ceph server.
- *CEPH-RGW-Port*—Port of the S3 gateway.
- *source-bucket-name*—the source bucket on which you want to enable bucket logging.

Expected output:

```
<BucketLoggingStatus
xmlns="http://s3.amazonaws.com/doc/2006-03-01/">

  <LoggingEnabled>

    <TargetBucket>log-destination</TargetBucket>

    <TargetPrefix>abc</TargetPrefix>

    <ObjectRollTime>300</ObjectRollTime>

    <LoggingType>Journal</LoggingType>

    <RecordsBatchSize>0</RecordsBatchSize>

    <TargetObjectKeyFormat>

      <PartitionedPrefix>

        <PartitionDateSource>EventTime</PartitionDateSource>

      </PartitionedPrefix>

    </TargetObjectKeyFormat>

  </LoggingEnabled>

</BucketLoggingStatus>
```

Creating policy for the log bucket

To create a log bucket policy using the Ceph console, refer to the IBM Storage Ceph documentation. This section describes how to create the policy using the S3 API.

Method: PUT**API endpoint:**

`http://<IP-of-CEPH-server>:<CEPH-RGW-Port>/<source-bucket-name>?logging`

Where:

- *IP-of-CEPH-server*—IP address of the Ceph server.
- *CEPH-RGW-Port*—Port of the S3 gateway.
- *source-bucket-name*—the source bucket on which you want to enable bucket logging.

API payload:

```
{  
  
  "Version": "2012-10-17",  
  
  "Statement": [  
  
    {  
  
      "Sid": "AllowLoggingFromSourceBucket",  
  
      "Effect": "Allow",  
  
      "Principal": {  
  
        "Service": "logging.s3.amazonaws.com"  
  
      },  
  
      "Action": "s3:PutObject",  
  
      "Resource": "arn:aws:s3:::<target-bucket-name>/*",  
  
      "Condition": {  
  
        "ArnLike": {  
  
          "aws:SourceArn": "arn:aws:s3:::<source-bucket-name>"  
  
        }  
  
      }  
  
    }  
  
  ]  
}
```

```

    }

]

}

```

Where:

- *<target-bucket-name>* is the bucket where target logs are generated.
- *<source-bucket-name>* is the bucket where you want to enable bucket logging.

Additional storage requirements at the staging location

You must reserve space in the backup host's staging location to store log files. See [“Configure a temporary staging location”](#) on page 17.

The required space is directly proportional to the number of tracked objects.

Use this formula to calculate an approximate estimate of the required additional storage space:

Estimated storage size (S) = 2 × Number_of_keys × Average_key_size (in bytes)

For example:

If the **Number_of_keys** = 1000000, and the **average_key_size** = 50 bytes

Estimated storage size (S) = 2 × 1000000 × 50 = 100,000,000 bytes = 100 MB

For details of maintaining the logs in Ceph, see: [Maintaining the log bucket](#).

Configuring bucket logging in IBM Storage Ceph

Follow these steps in your Ceph console for bucket logging configuration. Refer to IBM Storage Ceph documentation for details of any of these steps.

- 1 Create or identify the log bucket to store the log objects.
- 2 Enable the bucket logging on the source buckets with the following configuration. See [“Enable bucket logging for source buckets”](#) on page 54.
 - Each source bucket must have a unique prefix. It cannot be empty.
 - Logging type must be Journal. This logging type, logs only those operations that modify an object, such as PUT, POST, or DELETE.
 - Set the *TargetObjectKeyFormat* parameter with the value *PartitionedPrefix*.
 - Set the *PartitionDateSource* parameter as *EventTime*.

- For the parameter *ObjectRollTime* , specify the time in seconds after which a new log object is created. Default is 3600 seconds (one hour).
- 3 Configure a policy on the log bucket to enable logging for the other buckets. See [“Configuring access policy on the log bucket”](#) on page 52.
- 4 Configure access permissions on the source and log object buckets. See [“Configuring access permissions for the buckets”](#) on page 52.

Maintaining the log bucket

Log objects may accumulate in the log bucket depending on the source bucket's configuration and change rate. To manage storage and control costs, regularly delete log objects that are no longer needed. You can automate this process using Ceph's bucket lifecycle policies on the log bucket. To maintain the log bucket, you must configure at least one lifecycle policy.

The recommended policy is the S3 lifecycle policy on the log bucket, which offers the following benefits:

- Automatically deletes logs older than a specified number of days.
- Configurable retention period during policy creation.
- Validated by NetBackup against the backup schedule.
- Adjusts log pruning with full and differential backup frequency.

For details on configuring and using this lifecycle policy, refer to the IBM Storage Ceph documentation.

Configuring NetBackup for object change tracking

This section explains the configuration required to enable object change tracking for NetBackup.

- The primary server and backup host(s) must be of NetBackup version 11.1 or later.
- The credentials configured in NetBackup for the Cloud object store account must have the required permissions to access both the source and log buckets. See [“Configuring access permissions for the buckets”](#) on page 52.
- This feature is supported only when the NetBackup policy is configured with the Dynamic multi-streaming option.
- Object change tracking does not support a scale out server as a backup host.
- If you modify the object change tracking configuration, NetBackup verifies the updated settings. If the configuration is incompatible, the next backup fails. If

the modified configuration is compatible, NetBackup performs a full bucket scan during the next backup to detect object changes. The subsequent backups use object change tracking.

Configuring NetBackup policy for object change tracking

To enable your NetBackup policy to use object change tracking, select the option **Use object change tracking for incremental backup** in the **Cloud objects** tab while creating a backup policy. See [“Configuring the Cloud objects tab”](#) on page 78.

Verifying object change tracking in the Activity monitor

You can confirm if object change tracking has been activated for a policy in the Activity monitor. After a backup runs, open the Activity Monitor and check the logs. You should see the following log entry:

Backup using object change tracking, as specified in policy.

This message confirms that the backup is using the configured object change tracking mechanism.

Scenarios when NetBackup overrides object change tracking

In some cases, NetBackup may detect that object change tracking is insufficient to generate a list of changed objects due to system errors or manual disruptions. When this occurs, NetBackup switches to object scanning to identify changes since the last backup.

No bucket logging context

NetBackup stores a reference context for each backup in the log bucket to track the last backup. If this context is missing—due to deletion, pruning errors, or system issues—NetBackup cannot reliably identify changed objects. In such cases, NetBackup performs a full object scan.

Changes in object change tracking configuration

If you modify the object change tracking configuration, NetBackup verifies the updated settings. If the configuration is incompatible, the next backup fails. If the modified configuration is compatible, NetBackup performs a full bucket scan during the next backup to detect object changes. The subsequent backups use object change tracking.

Re-enable Object Change Tracking in NetBackup Policy

If you disable the **Use object change tracking for incremental backup** option, and then re-enable it in the policy, NetBackup performs the next backup by ignoring

object change tracking, and directly scanning the entire bucket to identify the changed objects. The subsequent backups are taken using object change tracking.

About storage lifecycle policies

A storage lifecycle policy (SLP) is a storage plan for a set of backups. Cloud object store policies let you use SLPs for optimized protection. 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 retrieves 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 use 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*.

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

Select **Backup**.

Destination storage attributes > Destination storage
Select the required storage.

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.

About policies for Cloud object store assets

Backup policies provide the instructions that NetBackup follows to back up objects. 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, 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 back up.

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 the 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 buckets or 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 several objects, while the other accounts are smaller with a 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.
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 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	Configure a temporary staging location	See “Configure a temporary staging location” on page 17.

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

Step	Action	Description
Step 5	Consider the backup schedule	Create additional backup policies if the schedules in one policy do not accommodate all the 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 do the objects change? If some objects change more frequently than others, or new objects get added in the bucket/container more frequently, 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 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.

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

Step	Action	Description
Step 6	Optimize performance with multi-streaming	NetBackup's performance in your environment depends on three major factors: ■ Network bandwidth between the backup host (media server) and the cloud storage service. ■ The cloud server's capability to handle multiple API requests. ■ The backup host's (media server) system memory (RAM). You must adjust the streams according to the number of objects that you are protecting, and your available system and network resources. NetBackup recommends 8–16 streams per policy. But you can specify the number of streams as per your environment.
Step 7	Select exactly what to back up.	You do not need to back up 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.
- In the **Cloud objects** tab, you must have permission to view and select the Cloud object store account(s), and the access host to specify the backup host or scale-out server for a policy.
- Evaluate the requirement for NetBackup accelerators in your environment. If you want to use accelerators, you need to specify this while creating the policy.
- If you plan to use a backup host or scale-out server apart from the one used to validate your cloud object store account, ensure that the necessary ports are open and configurations are in place. This is crucial for enabling server communication with the cloud provider's endpoint through REST API calls.

You can use a scale-out server if you have a large number of buckets in your Cloud object store. NetBackup Snapshot Manager can scale out as many data mover containers as needed at run time, and then scale them down when the data protection jobs are completed. You do not need to worry about configuring multiple backup hosts, and creating multiple policies to distribute the load across these backup hosts.

- 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 this while creating the policy. To use multistream, you also need to configure the number of jobs for the buckets as clients 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 "Policy attributes" on page 68.
Schedule your backups.	See "Creating schedule attributes for policies" on page 72. See "Configuring the Start window" on page 74. See "Configuring the exclude dates" on page 76. See "Configuring the include dates" on page 77.
Select the account and objects to backup.	See "Configuring the Cloud objects tab" on page 78. See "Adding conditions" on page 80. See "Adding tag conditions" on page 81. See "About dynamic multi-streaming" on page 49.

Policy attributes

The following procedure describes how to select the attributes for the backup policy.

Select the policy attributes

- 1 On the left, click **Protection > Policies**.
- 2 Enter a name for the policy in the **Policy name** field.
- 3 Select the **Cloud-Object-Store** option from the **Policy type** dropdown.
- 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 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** dropdown. 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.
- 5 **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. A retry is often quicker than restarting the entire job.

The checkpoint frequency indicates how often NetBackup takes a checkpoint during a backup. The default is 30 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.

- 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 **Limit jobs per policy** 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 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.
- 9 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 it. 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 are available in the **Policies** list.
- 10 The **Allow multiple data stream** option is selected by default and is read-only. This option allows NetBackup to divide automatic backups for each query into multiple jobs. Because the jobs are in separate data streams, they can occur concurrently.

Multi-stream jobs consist of a parent job to perform stream discovery and child jobs for each stream. Each child job displays its 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.

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

NetBackup Accelerator optimizes backups to perform better both in terms of data movement and backup time. Accelerator identifies changed content, move only unique data to backup target. Subsequently, it creates a synthetic backup image by combining current changed data and unchanged data from the previous backup image. 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 backup data that is stored in previous backups.

Typically, any data change for an object results in an update of its modification time (mtime). Certain Cloud object store vendors may not update the mtime for an object's tag or user attribute change. NetBackup's Cloud object store component considers object's modification time (mtime) to identify the objects that have changed since the last backup. Additionally, it must check if tags or user attributes has changed as well to account for the cloud vendor's inability to update mtime for tags change.

Optionally, select the **Quick object change scan** option to skip the checks for object tags. If you have an application environment, where the object tags are not modified after the initial create time, you can use this option to skip these checks and significantly increase the backup speed.

Enabling this option speeds up object change identification for Accelerator, as it skips the comparison of object tags since last backup. However, with this option you may not backup objects that have no data change but only tag change since the last backup.

This feature works only with Dynamic multi-streaming. You need to configure a temporary storage space for backups before using this option. See [“Configure a temporary staging location”](#) on page 17.

- 12 Select the **Disable for all clients** option from the Client-side deduplication options. NetBackup Cloud object store protection uses the backup host as the client.
- 13 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 the **Type of backup**:
 - **Full Backup**-A complete backup of the objects that contain all of the data objects and the log(s).
 - **Differential Incremental Backup**-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 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 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** to 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 periods during which NetBackup can start jobs when using a schedule. Periods are referred to as windows. Configure the windows to 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.

Adding, changing, or deleting a time window in a policy schedule

Use one of the following procedures to add, change, or delete a time window.

To configure the Start window:

- 1 On the left, click **Policies**, under **Protection**. 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'd like the window to start and drag it to the day and time when you'd 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'd like the window to start and drag it to the day and time when you'd 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 buckets/containers with backups due.



The diagram illustrates the following points:

Point 1

The bucket/container A3 starts within the Schedule A time window but does not complete until after the Schedule B start time. However, the bucket/container A3 runs to completion even if the window closes while the backup is running. Bucket/container B1, on Schedule B, begins as soon as bucket/container A3 completes.

Point 2

Schedule A does not leave enough time for all the bucket/containers on Schedule B to be backed up. Consequently, the bucket/container B3 is unable to start because the time window has closed. The bucket/container B3 must wait until the next time NetBackup runs Schedule B.

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 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 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.
 - 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 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. 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 objects in the desired buckets.

NetBackup lets you make a 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 a single backup host or scale-out server 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 or scale-out servers.

To configure cloud objects:

- 1 Select a **Cloud object store account** and **Host**. You can see a list of accounts and backup hosts that you are privileged to access. If you use a scale-out server for the account, the **Host** field is disabled. You cannot change a scale-out server while creating a policy.
- 2 (Optional) Select the **Allow dynamic multi-streaming** option to allow NetBackup to divide automatic backups for each bucket/container into concurrent multiple streams. This option can dramatically improve the backup time of the protected buckets/containers.

In the field, **Maximum number of streams per bucket/container** specify a number between 1 to 64. The default value is 8.

Some streams may go to a queue, if the maximum number of concurrent jobs allowed in the storage unit selected for the policy, is less than the total number of streams that are running for the policy. For optimal performance, keep the **Maximum concurrent jobs** allowed property of selected storage greater than the total number of streams that you expect the policy to handle. The minimum value for the **Maximum concurrent jobs** for the storage must be 64.

Optionally, specify a **Temporary staging location** path, or use the default path. You can leave this field blank to use the default path. See [“Configure a temporary staging location”](#) on page 17.

Note: If you enable dynamic multi-streaming, all selected buckets and containers are completely backed up. You cannot define any queries for the buckets or containers that you have selected.

If you select Ceph as the service provider, you can enable bucket logging by selecting **Use object change tracking for incremental backups**. Note that this disables the **Quick object change scan** option in the **Attributes** tab.

- 3 To add buckets/containers, click **Add** near 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/container 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.

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.

- 4 To add a query to the selected buckets/containers, click **Add query** under **Queries**.
- 5 Enter a name for the query, and select the buckets that you want to filter using the query.
- 6 In the **Select objects/blobs** table, select the option **Include all objects/blobs in the selected buckets/containers** to backup one or more entire buckets.
- 7 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 80. and See [“Adding tag conditions”](#) on page 81. respectively, for more details.

Adding conditions

NetBackup gives you the convenience of selectively backing up 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.

If you enable dynamic multi-streaming, all selected buckets and containers are completely backed up. You cannot define any queries for the buckets or containers that you have selected.

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 the query, and 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 objects.

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 the **Query name** and **Queries** columns. The values of the **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, and 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.

This feature is not available for Azure Data Lake Storage and Azure Data Lake Storage Government providers.

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 the query, and 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 objects.

- 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 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.

Examples 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:

- The following blobs are tagged with "Project": "HR" tag
 - OrganizationData/Hr/resumes/resume1_selected.pdf
 - OrganizationData/Hr/resumes/resume2_rejected.pdf
 - OrganizationData/Hr/resumes/resume3_noupdate.pdf
- The following blobs are 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
- The following blobs are 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
- The following blobs are 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 back up all blobs tagged with "Project": "Finance", add the following query:

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

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

- Case 2: To back up 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 back up 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 backups for a policy.

View Cloud object store policies

- 1 On the left, click **Policies**. All the policies that you have the 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 68.

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.

Recovering Cloud object store assets

This chapter includes the following topics:

- [Prerequisites for recovering Cloud object store objects](#)
- [Configuring Cloud object retention properties](#)
- [Recovering Cloud object store assets](#)

Prerequisites for recovering Cloud object store objects

Before you start recovery, make sure that these conditions are met.

- Keep handy information about the destination bucket(s) or container(s) that you want to use for recovery.
- Ensure that you set the **Client connection timeout** and **Client read timeout** parameters to 3600 seconds, in the primary server. You can set these parameters in **Host properties**.
 - On the left, click **Hosts**, and then click **Host properties**.
 - Select the primary server, and click **Edit primary server**.
 - Click **Timeout** on left, and enter value of the parameters **Client connection timeout** and **Client read timeout** as 3600 seconds. Click **Save**.
See the *NetBackup Web UI Administrator's Guide* for more details.
- Select the object(s) that you want to recover. You can recover objects by selecting all objects or blobs from the selected image. Alternatively, select

individual objects, select all objects under a set of folder(s), or all objects matching a set of prefixes.

- A valid Cloud object store account is required to access the buckets, containers, 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 the ones required for backup. If it helps, you can create a 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 a 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 or scale-out server 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 a folder or prefix.

Configuring Cloud object retention properties

Starting from NetBackup 10.3, the object lock feature lets you retain the original object lock properties and also provides an option to customize the object lock properties. When you apply object lock properties on restored objects, the restored objects cannot be deleted until the retention period is over, or legal holds are removed. No configuration is needed during policy creation and backup to use the object lock and retention properties backup.

Note: This option can incur additional cloud storage costs to hold data for a longer time. Avoid using these options if you want a temporary copy of data that you want to delete after browsing the objects or copying to another location.

To apply object retention locks or legal holds on the restored objects, you can select multiple options during restoration to meet your organization's compliance and retention requirements. You can select the options in the **Recovery options** page, under **Advanced restore options**. See [“Recovering Cloud object store assets”](#) on page 87.

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 with completed backups. You can search the table using the search box.
 - To add a specific bucket or container, select **Add the bucket/container details** option. If you have selected an Azure Data Lake workload, select **Add files/directories**.
Select the cloud provider, and enter the bucket/container name, and the Cloud object store account name. For Azure workloads, specify the storage account name, if available in the UI.

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 the catalog view as a backup ID. You can select the bucket by manually entering the bucket name, provider ID, and the Cloud object store account name as per the backup ID. The backup ID is formed as

`<providerId>_<cloudAccountname>_<uniquename>_<timestamp>`

for Azure the `uniquename` is `storageaccountname.bucketname`, and for S3 providers it is the bucket name.

- 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 In the Recovery details page, you can add the objects and folders or prefix and scan the selected images for malware before restoring the images:
 - (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. For an Azure Data Lake workload, this option is available as **Include all files/directories**. You can use the left navigation tree structure to filter the table. Click **Add**.

The following warning message is displayed when images which are not scanned are selected for recovery:

```
One or more images selected for recovery are not scanned.
```

Note: To restore from malware-affected images, you must have the Administrator role or equivalent RBAC permissions.

For more information on recovering from malware infected images, see *Security and Encryption Guide*.

- (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**.

Note: Clean file recovery (Skip infected files) as part of recovery is not supported for Cloud-Object-Store.

8 In the Recovery options page, you can select whether you want to restore to the source bucket of the container or use a different one. 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.
 - If you have selected an Azure Data Lake workload, enter the **Directory to restore**.
- **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 the 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 create accounts with limited and specific permissions for backup and restore. In this case, you can provide the same bucket as the original to restore to the 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 the destination from the list.
 - 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.

- 9 Select a **Recovery host**. The recovery host that is associated with the Cloud object store account is displayed by default. If required, change the **Backup host**. If the Cloud object store account uses a scale-out server, this field is disabled.
- 10 Optionally, to overwrite any existing object or blobs using the recovered assets, select **Overwrite existing objects/blobs**.
- 11 (Optional) To override the default priority of the restore job, select **Override default priority**, and assign the required value.
- 12 In the **Advanced restore options**:
 - To apply the original object lock attributes from the backed-up objects, select **Retain original object lock properties**.
 - To change the values of different properties, select **Customize object lock properties**. From the **Object lock mode** list:
 - Select **Compliance** or **Governance** for Amazon or other S3 workloads.
 - Select **Locked** or **Unlocked** for Azure workloads.
 - Select a future date and time till which the object lock is valid. Note that the recovered object is locked till this specified date and time.
 - Select **Object lock legal hold status** to implement it on the restored objects.

See [“Configuring Cloud object retention properties”](#) on page 87.

The **Advanced restore options** are not applicable to the Azure Data Lake workload.

- 13** In the Review page, view the summary of all the selections that you made, and click:

- **Start recovery**

Or

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

Troubleshooting

This chapter includes the following topics:

- Credential validation failed. Provide correct Cloud account name and credentials.
- Error 5549: Cannot validate bucket logging information
- Error 5576: The maximum number of concurrent jobs specified for a storage unit, must be greater than or equal to the number of streams specified in the policy.
- Error 5579: Falling back to object listing for change detection, not considering object change tracking for this bucket, specified in the policy.
- Error 5580: The specified failover strategy for the storage unit group is incompatible with the Cloud object store policy, with dynamic multi-streaming.
- Error 5545: Backup failed as NetBackup cannot parse records from the log object
- Error 5541: Cannot take backup, the specified staging location does not have enough space
- Error 5537: Backup failed: Incorrect read/write permissions are specified for the download staging path.
- Error 5538: Cannot perform backup. Incorrect ownership is specified for the download staging path.
- Reduced acceleration during the first full backup, after upgrade to versions 10.5 and 11.
- After backup, some files in the shm folder and shared memory are not cleaned up.
- After an upgrade to NetBackup version 10.5, copying, activating, and deactivating policies may fail for older policies

- Backup fails with default number of streams with the error: Failed to start NetBackup COSP process.
- Backup fails, after you select a scale out server or Snapshot Manager as a backup host
- Backup fails or becomes partially successful on GCP storage for objects with content encoded as GZIP.
- Recovery for the original bucket recovery option starts, but the job fails with error 3601
- Recovery Job does not start
- Restore fails: "Error bpbrm (PID=3899) client restore EXIT STATUS 40: network connection broken"
- Access tier property not restored after overwriting the existing object in the original location
- Reduced accelerator optimization in Azure for OR query with multiple tags
- Backup failed and shows a certificate error with Amazon S3 bucket names containing dots (.)
- Azure backup jobs fail when space is provided in a tag query for either tag key name or value.
- The Cloud object store account has encountered an error
- The bucket is list empty during policy selection
- Creating a second account on Cloudian fails by selecting an existing region
- Restore failed with 2825 incomplete restore operation
- Bucket listing of a cloud provider fails when adding a bucket in the Cloud objects tab
- AIR import image restore fails on the target domain if the Cloud store account is not added to the target domain
- Backup for Azure Data Lake fails when a back-level media server is used with backup host or storage server version 10.3
- Backup fails partially in Azure Data Lake: "Error nbpem (pid=16018) backup of client
- Recovery for Azure Data Lake fails: "This operation is not permitted as the path is too deep"

Credential validation failed. Provide correct Cloud account name and credentials.

- Empty directories are not backed up in Azure Data Lake
- Recovery error: "Invalid alternate directory location. You must specify a string with length less than 1025 valid characters"
- Recovery error: "Invalid parameter specified"
- Restore fails: "Cannot perform the COSP operation, skipping the object: [/testdata/FxtZMidEdTK]"
- Cloud store account creation fails with incorrect credentials
- Discovery failures due to improper permissions
- Restore failures due to object lock

Credential validation failed. Provide correct Cloud account name and credentials.

Explanation

This issue may occur in the following scenarios:

- The system is running RHEL 9.6 or later.
- The OpenSSL version was manually upgraded.

In newer RHEL releases, OpenSSL has been updated from version 3.2.2 to 3.5.1. The newer version has known issues for NetBackup functionality. As a result, both discovery and backup might fail in NetBackup.

To verify the issue, you can check the `nbcosp_sys` log file on the backup host, at the location:

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

Look for the entry:

```
panic: tls: HKDF-Extract invocation failed unexpectedly:
EVP_PKEY_CTX_set1_hkdf_salt
    openssl error(s):
    error:078C0102:common libcrypto routines::passed a null
parameter
```

Workaround

1. Add the following entries to the `bp.conf` file on the backup host:
 - `COSP_DISABLEKYBER = 1`

- `COSP_FORCETLSV12 = 1`
- 2. Restart the `nbcosp` service.
- 3. Retry the operation.

Error 5549: Cannot validate bucket logging information

Explanation:

This may occur due to the following reasons:

- Validation failed because the partition prefix is either missing or incorrectly set. Configure bucket logging with the correct date partition prefix.
- Cannot retrieve the target bucket.
- Unsupported `LoggingType` attribute.
- Missing or invalid target prefix.
- Cannot retrieve the `PartitionDateSource` attribute.

Workaround:

Check the configuration and retry the backup operation. Configurations should be like:

- Each source bucket must have a unique prefix. The prefix cannot be empty.
- The `LoggingType` attribute must be *Journal*.
- Set the `TargetObjectKeyFormat` attribute as *PartitionedPrefix*.
- Set the `PartitionDateSource` attribute as *EventTime*.

Error 5576: The maximum number of concurrent jobs specified for a storage unit, must be greater than or equal to the number of streams specified in the policy.

Explanation:

This error occurs when the concurrent job limit set for a storage unit is lower than the number of streams specified in the policy.

Workaround:

Error 5579: Falling back to object listing for change detection, not considering object change tracking for this bucket, specified in the policy.

Increase the concurrent jobs setting for the storage unit to match the number of streams specified in the policy.

Error 5579: Falling back to object listing for change detection, not considering object change tracking for this bucket, specified in the policy.

Explanation:

This issue may occur due to the following reasons:

- The marker is empty or invalid. This can happen if the previous backup image is missing or if this is the first full backup.
- The stream count has changed since the last backup.
- The bucket logging configuration has changed.
- The `tracklog` from the previous backup is missing.
- A forced rescan is selected in the schedule.

Workaround:

There is no workaround for this issue. For all these reasons, NetBackup uses object listing for change detection. NetBackup will revert to object change tracking for the subsequent backups.

Error 5580: The specified failover strategy for the storage unit group is incompatible with the Cloud object store policy, with dynamic multi-streaming.

Explanation:

This occurs when the storage unit group is incompatible with the Cloud object store policy.

Workaround:

Error 5545: Backup failed as NetBackup cannot parse records from the log object**To change storage unit selection:**

- 1 In the NetBackup UI, on the left, click **Storage**, then click **Storage units**. Click the **Storage unit group** tab.
- 2 Click **Add** to add a new storage unit, or edit an existing one.
- 3 Under **Storage unit selection** options, select any of these options as required: **Prioritized**, **Round Robin**, or **Media Server Load Balancing**.

Retry the backup.

Error 5545: Backup failed as NetBackup cannot parse records from the log object

Explanation:

When the **Use object change tracking for incremental backups** feature is enabled in the policy, the backup operation might fail due to a parsing error. This error occurs while processing the log file, which contains the names of changed objects along with other fields.

The error can occur if object names contain unsupported or hard-to-handle characters, such as newline characters.

The `nbcosp` log displays the following message:

```
{"level":"error","JobId":167,"Failed to parse record from logging
object":"not enough fields after
uation","time":"2025-10-16T12:13:10.11664034+05:30","caller":"main.parseBackupLog.nbcosp_request_handler.go:261","message":"Parsing
object line"}
```

Workaround:

Ensure that object names do not contain unsupported characters, such as newline characters.

Error 5541: Cannot take backup, the specified staging location does not have enough space

Explanation:

The staging location does not have sufficient storage space required to download the Cloud object store object data, causing the backup to fail.

Workaround:

Increase the storage space and try the backup again.

Error 5537: Backup failed: Incorrect read/write permissions are specified for the download staging path.

Error 5537: Backup failed: Incorrect read/write permissions are specified for the download staging path.

Explanation:

This error occurs when the NetBackup service user on the backup host does not have read and write permission of the cache DB path.

Workaround:

Provide the read and write permissions to the NetBackup service user. Try the backup again.

Error 5538: Cannot perform backup. Incorrect ownership is specified for the download staging path.

Explanation:

This error occurs if the NetBackup service user does not have ownership of the directory of the cache DB path on the backup host.

Workaround:

Provide ownership of the directory of the cache DB path on the backup host to the NetBackup service user. Try the backup again.

Reduced acceleration during the first full backup, after upgrade to versions 10.5 and 11.

Explanation:

In NetBackup versions 10.5 and 11, we have changed some metadata, which gets protected as part of the backup across releases. To track the newer metadata format, all objects are considered changed due to a mismatch in the metadata properties of the blob.

Workaround:

This is expected behavior.

After an upgrade, NetBackup does a full rescan backup to avoid dependency of images across releases.

After backup, some files in the `shm` folder and shared memory are not cleaned up.

Explanation:

In backups with dynamic multi-streaming, a crash of some NetBackup process may cause the backup to be partially successful or fail. In such cases, some files and the shared memory are not cleaned up. These files and buffers are small and should not take up much space.

Workaround:

Do the following:

- To clear files, go to the folder on the backup host:
`<install_directory>/netbackup/db/config/shm/`
 These folders contain files in the format: `<parentid_streamNumber>`, for the failed or partial backups that have finished. Delete the files.
- To clear the shared memory:
 - You can view the shared memory in the backup host, using the command:
`ipcs -m`
 - Wait for all the running backups and restores to finish, and then restart the backup host. This clears the shared memory.

After an upgrade to NetBackup version 10.5, copying, activating, and deactivating policies may fail for older policies

Explanation:

You can see the message:

The attribute 'useMultipleDataStreams' is false. You must enable this attribute.

NetBackup version 10.5, makes it mandatory for the policies to use multiple data streams, whereas the older policies do not have this attribute.

Workaround:

1. Edit the older policy, check if the allow multiple data streams option is selected by default. Save the policy.
2. Retry the operation.

Backup fails with default number of streams with the error: Failed to start NetBackup COSP process.

Explanation:

This error occurs because the `nbcosp` service, responsible for making calls to the cloud provider's APIs, is down or crashed.

Workaround:

Do any of the following:

- On the backup server, run the following command to start the `nbcosp` service.
`<install_directory>/pdde/pdcr/bin/nbcosp start`
- On the backup server, start all the NetBackup services, run:
`<install_directory>/netbackup/bin/bp.start_all`

Backup fails, after you select a scale out server or Snapshot Manager as a backup host

Explanation:

The error occurs when a policy has Dynamic multi-streaming, and you update the Cloud object store account used in the policy to use a scale-out server or NetBackup Snapshot Manager as a backup host.

Workaround:

Do any of the following:

- Disable Dynamic multi-streaming in the policy.
- Do not use a scale out server or NetBackup Snapshot Manager as a backup host in the Cloud object store account.

Backup fails or becomes partially successful on GCP storage for objects with content encoded as GZIP.

Explanation:

Recovery for the original bucket recovery option starts, but the job fails with error 3601

The problem is attributed to the 'decompressive transcoding' feature of GCP. When an object is compressed with gzip and the content-encoding header is set to gzip, the objects go under decompressive transcoding and they are stored compressed in cloud. Such objects never report object size as part of the content length header. NetBackup cannot backup an object when the size is not returned.

For details, see this [GCP documentation](#) page.

Workaround:

Use a different mode of content encoding than GZIP.

Recovery for the original bucket recovery option starts, but the job fails with error 3601

Explanation

This error occurs because any of the four reasons:

- Cloud object store account required to connect to the cloud to do recovery does not exist.
- The Cloud object store account used during the backup of the bucket does not exist in the 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 the 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". Occurs even if the Cloud object store account with the same name is used during backup.

Workaround

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

Do the following:

- 1 Use the same Cloud object store account in the 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 the 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 a network error message in the Activity monitor, change the system configuration time-out to 900/1200 or the appropriate high value, in seconds. Start a new restore job.

Steps to set a time-out are as follows: On the left, click **Hosts > Host properties** > Select **Primary/ Media server** > select the **Time-out** option > set the time-out value -> click **Save**. Refer to the *Web UI Administrator Guide* for details.

Access tier property not restored after overwriting the existing object in the 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 the case of Azure cloud storage, when we have an object with accessTier as cool, and we try to upload an object/blob of the same name with Hot(inferred) accessTier with an overwrite option, accessTier remains as cold. A new access tier

does not get set. This behavior is observed when a file is uploaded from the portal. It does not change accessTier from cool to Hot (inferred) when the Overwrite option is selected on the Azure portal.

Reduced accelerator optimization in Azure for OR query with multiple tags

Explanation

When the Cloud object store policy has at least one query with multiple tag conditions, combined using the "OR" operator, the backups of Azure containers using accelerator enabled Cloud object store policy, it shows a 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 the tracklog; hence, such objects are backed up repeatedly without getting accelerator benefit for these objects.

Workaround

Do not use the OR condition while combining multiple tag conditions for Azure. Instead, create a 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 the first backup, which is without any acceleration for these new queries. For subsequent backups, you can see the problem is resolved.

Backup 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 the bucket:` Since the path-style URL adds the bucket as a part of the URL path and not as a host name, we do not get any

Azure backup jobs fail when space is provided in a tag query for either tag key name or value.

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 an older S3 URL as a path style and can connect with a bucket with a (.) in the name. To do this, you can add a region with a plain S3 endpoint (`s3.<region-id>.amazonaws.com`) and select the URL Access Style as the path style.

- Disable SSL: This workaround is not the recommended one, since it replaces the secure endpoint with an unsecure/unencrypted endpoint. After turning off SSL, it disables the peer-host validation of the server certificate. It bypasses the host name match for the virtual host-style URL of bucket (`bucket.123.s3.dualstack.us-east-1.amazonaws.com`) with the subject name in the certificate (`*.s3.dualstack.us-east-1.amazonaws.com`).

Azure backup jobs fail when space is provided in a tag query for either tag key name or value.

Workaround

Do not use spaces in the 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 an error state when:

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

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

Workaround

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

Case 1: When there is no valid Cloud storage server or MSDP-C LSU with the same name as the 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 the csconfig CLI and note down the details of the alias.
 - Use the 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 the following command to get the details of the 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 the 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 that is in an error state.
- Create the Cloud object store account using the noted details.

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

- You cannot reuse 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 an 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.

The bucket is list empty during policy selection

Explanation:

In NetBackup, you can add 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 clouds might not have a region configured.

When you are using such an invalid region in an account, the bucket list 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 the region location constraint.
- 2 Correct the region details by editing the account configuration. See [“Adding Cloud object store accounts”](#) on page 29.
- 3 To edit the 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 a second account on Cloudian fails by selecting an existing region

Explanation:

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

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

Workaround:

Use the NetBackup Asset Query APIs to create an account. The region details part can be provided in the 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 the 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 are restored from the backup image. Restore failed with 2825 incomplete restore operation.

Explanation:

This error can occur for multiple reasons. The most likely scenario for this error is when a cloud API initiated by NetBackup during a restore returns an error like an HTTP 400 status code (Bad Request). The reasons can vary with each cloud vendor. For example, GCP supports different Content-Language metadata as compared

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

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` log shows the following message:

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

`nbtar` log shows the following message:

```
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 settings.

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 the 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 a cloud provider fails when adding a bucket in the 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.

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

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 implies that no request can be issued to the cloud provider without providing a bucket name as the 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 an 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`, and `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 the `/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 to the target domain

Error

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

Explanation

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

Workaround**Solution 1:**

Create the Cloud object store account in the target domain with the same name as in the source domain and perform restore. See [“Adding Cloud object store accounts”](#) on page 29.

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 the source account name. Click **Next**.
- 2 Select the backup image and add objects, 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 restoration.

Backup for Azure Data Lake fails when a back-level media server is used with backup host or storage server version 10.3

Error message

bpbkar Exit: INF - EXIT STATUS 3600: Cannot perform the COSP operation.

Error nbpem (pid=13052) backup of client
azuredatalake_COSv17_adlsgen2xxxxx.xxxxxxx exited with status 3600 (cannot perform the COSP operation).

Explanation

Support for Azure Data Lake was introduced in NetBackup version 10.3, so this workload does not work on media or storage server versions before 10.3.

Workaround

Ensure that media server, backup host or scale-out server, and storage server versions are 10.3 or later.

Backup fails partially in Azure Data Lake: "Error nbpem (pid=16018) backup of client

Explanation

Happens if you have a directory name + file name (including all '/') longer than or equal to 1,024 characters.

Workaround

Recovery for Azure Data Lake fails: "This operation is not permitted as the path is too deep"

You must have a path consisting of directory path + file name (including all '/') under 1,024 characters. The maximum directory length allowed is 1023 characters, including the "/" in the beginning and the end of the path.

Recovery for Azure Data Lake fails: "This operation is not permitted as the path is too deep"

Explanation

Happens when the selected directory depth for recovery is more than 60. You can see this error message in the `nbcosp` logs.

Workaround

You must use a path with directory depth less than 60, excluding container.

Empty directories are not backed up in Azure Data Lake

Explanation

Occurs when you add an empty directory or a leaf level directory without any contents, to a policy's query filter. During backup that empty directory is not backed up.

Workaround

Select the **Include all files/directories in the selected buckets/containers** option to backup the empty directories.

Recovery error: "Invalid alternate directory location. You must specify a string with length less than 1025 valid characters"

Explanation

Happens when the specified value in the **Directory to restore** field in the web UI, is more than 1024 characters. This value maps internally to the **alternateDirectoryLocation** attribute in the API.

Workaround

Specify a restore location that is less than 1024 characters.

Note: This error does not appear on the **Activity Monitor**.

Recovery error: "Invalid parameter specified"

Explanation

Happens when the API contains values for **alternateDirectoryLocation** and **addPrefix** attributes. The API does not support one of the specified input parameters.

Workaround

Azure Data Lake Storage Gen2 supports only the **AlternateDirectoryLocation** attribute and thus the **addPrefix** attribute must be empty in the API request.

Note: This error does not appear on the **Activity Monitor**. For more details, refer to NetBackup API documentation.

Restore fails: "Cannot perform the COSP operation, skipping the object: [/testdata/FxtZMidEdTK]"

Explanation

Restore jobs fail when the files or directories that are uploaded to Azure Data Lake Gen2 account using DFS API's Overwrite function, replacing the original files.

In the **nbcosp** logs as you can see the error:

```
*RESPONSE ERROR (ServiceCode=BlobOperationNotSupported)
====\nDescription=Blob operation is not supported.
```

Note: In this issue, files that are uploaded to Azure portal using the **Upload** option are not included.

Workaround

Do any of the following:

- Try restore operation on the same container to a different directory.
- Try restore operation to a different container.
- Try restore operation to a different destination.

- Delete the original directory, and then try to restore to the same location.

Cloud store account creation fails with incorrect credentials

For AWS accounts:

Viewing the `nbcosp` logs.

Location:

- For backup host: `/usr/opensv/netbackup/logs/nbcosp`
- For scale-out server:
`/cloudpoint/opensv/dm/datamover.<datamover_id>/netbackup/logs/nbcosp`

```
{ "level": "error", "Error Code": "SignatureDoesNotMatch", "Message": "The
request signature we calculated does not match the signature you
provided. Check your key and signing method.",
"level": "error", "errmsg": "Unable to validate creds.", "storage
server": "aws-acc", "time": "2023-07-
```

Viewing the `ncfnbcs` logs.

Location:

- For backup host: `/usr/opensv/logs/ncfnbcs`
- For scale-out server:
`/cloudpoint/opensv/dm/datamover.<datamover_id>/logs/ncfnbcs`

```
2,51216,309,366,474,1690282728130,1673,140536982484736,0:,0:,0:,2,(28|S113:ERR
- OCSD reply with error,error_code=1003 error_msg:
updateStorageConfig Failed as credential validation failed|)
2,51216,309,366,475,1690282728131,1673,140536982484736,0:,0:,0:,2,(28|S60:ERR
- operation_to_ocsd failed, storageid=aws-acc, retval=23|)
0,51216,526,366,6,1690282728131,1673,140536982484736,0:,132:Credential
validation failed for given account,
```

Workaround:

Update the credentials and try to create the account again.

Discovery failures due to improper permissions

Viewing the `nbcosp` logs.

Location:

- For backup host: `/usr/opensv/netbackup/logs/nbcosp`
- For scale-out server:
`/cloudpoint/opensv/dm/datamover.<datamover_id>/netbackup/logs/nbcosp`

```
25T11:14:14.761525555Z", "caller": "main.(*OCSS3).
listBucketsDetailsCOSP:s3_ops.go:5261", "message": "Unable
to listBucketsDetailsCOSP" } {"level": "debug", "status code"
:403, "errmsg": "AccessDenied: Access Denied\n\tstatus code: 403,
request id: K7JVVPWAGW4KYSQ6, host id:
```

Workaround:

Add the required permissions. See [“Permissions required for Amazon S3 cloud provider user”](#) on page 21.

Restore failures due to object lock

Explanation:

During restore, if you select the **Retain original object lock properties** option, NetBackup applies the object lock properties.

Viewing the Activity monitor logs:

```
Warning bpbrm (pid=21103) from client
ip-10-176-97-167.us-east-2.compute.internal: WRN - Cannot set Object
lock on the object. Access to perform the operation was denied.
```

```
Jul 25, 2023 11:26:00 AM - Error bpbrm (pid=21103) from client
ip-10-176-97-167.us-east-2.compute.internal: ERR - Cannot complete
restore for any of the objects.
```

```
Jul 25, 2023 11:26:00 AM - Warning bpbrm (pid=21103) from client
ip-10-176-97-167.us-east-2.compute.internal: WRN - The 3 files
restored partially as object lock cannot be applied.
```

```
Jul 25, 2023 11:26:00 AM - Info tar (pid=1697) done. status 5
```

Viewing the `nbcosp` logs:

```
{
  "level": "info",
  "SDK log body": "<?xml version='1.0' encoding='UTF-8'>\n<Error><Code>AccessDenied\n\n</Code><Message>Access Denied\n\n</Message><RequestId>ZNT4GXHP70HX573A</RequestId>\n\n</HostId>\n\n3scBnke9LmOwtuK5lnYv0ozyKqone+ey04qXtSt6s/OQbpSCyfxiwwdi2CPG3cHUH/z7C3mHeoX5Cnrb2xg=</HostId>\n\n</Error>\n",
  "time": "2023-07-25T05:56:00.708117368Z",
  "caller": "internal/logging.ExtendedLog.Log:zerolog_wrapper.go:18",
  "message": "SDK log entry"
}
{
  "level": "debug",
  "status code": 403,
  "errmsg": "AccessDenied: Access Denied\n\n\tstatus code: 403, request id: ZNT4GXHP70HX573A, host id: 3scBnke9LmOwtuK5lnYv0ozyKqone+ey04qXtSt6s/OQbpSCyfxiwwdi2CPG3cHUH/z7C3mHeoX5Cnrb2xg=",
  "time": "2023-07-25T05:56:00.708145345Z",
  "caller": "main.s3StatusCode:s3_ops.go:8447",
  "message": "s3StatusCode(): get http status code"
}
{
  "level": "error",
  "error": "AccessDenied: Access Denied\n\n\tstatus code: 403, request id: ZNT4GXHP70HX573A, host id: 3scBnke9LmOwtuK5lnYv0ozyKqone+ey04qXtSt6s/OQbpSCyfxiwwdi2CPG3cHUH/z7C3mHeoX5Cnrb2xg=",
  "time": "2023-07-25T05:56:00.708160142Z",
  "caller": "main.(*OCSS3).commitBlockList:s3_ops.go:2655",
  "message": "s3Storage.svc.PutObjectRetention Failed to Put ObjectRetention"
}
```

Workaround:

You must have the required permissions for object retention. These are the necessary permissions that your role must have:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ObjectLock",
      "Effect": "Allow",
      "Action": [
        "s3:PutObjectRetention",
        "s3:GetObjectRetention",
        "s3:DeleteObjectRetention"
      ],
      "Resource": "*"
    }
  ]
}
```

```

    "Action": [
        "s3:PutObjectRetention",
        "s3:BypassGovernanceRetention"
    ],
    "Resource": [
        "*"
    ]
}
]
}

```

See [“Configuring Cloud object retention properties”](#) on page 87.