

NetBackup™ Cloud Administrator's Guide

UNIX, Windows, Linux

Release 11.2.0.1

NetBackup™ Cloud Administrator's Guide

Last updated: 2026-09-07

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	About NetBackup cloud storage	8
	About cloud storage features and functionality	8
	About the catalog backup of cloud configuration files	11
	About support limitations for NetBackup cloud storage	12
Chapter 2	About cloud storage	14
	About the cloud storage vendors for NetBackup	14
	About the Amazon S3 cloud storage API type	15
	Amazon S3 cloud storage vendors certified for NetBackup	16
	Amazon S3 storage type requirements	16
	Permissions required for Amazon S3 cloud provider user	18
	Amazon S3 cloud storage provider options	18
	Amazon S3 cloud storage options	21
	Amazon S3 cloud storage server configuration options	23
	Amazon S3 credentials broker details	26
	About private clouds from Amazon S3-compatible cloud providers	27
	About Amazon S3 storage classes	28
	Amazon virtual private cloud support with NetBackup	29
	About protecting data in Amazon for long-term retention	31
	Protecting data using Amazon S3 Intelligent Tiering (LIFECYCLE)	42
	About using Amazon IAM roles with NetBackup	45
	About NetBackup character restrictions for Amazon S3 cloud connector	48
	Protecting data with Amazon Snowball and Amazon Snowball Edge	50
	About Microsoft Azure cloud storage API type	62
	Microsoft Azure cloud storage vendors certified for NetBackup	63
	Microsoft Azure storage type requirements	63
	Microsoft Azure cloud storage provider options	64
	Microsoft Azure advanced server configuration options	66
	Protecting data in Microsoft Azure Archive for long-term retention	68

Chapter 3	Configuring cloud storage in NetBackup	70
	Before you begin to configure cloud storage in NetBackup	71
	Configuring cloud storage in NetBackup	72
	Cloud installation requirements	73
	Scalable Storage properties	74
	Configuring advanced bandwidth throttling settings	75
	Advanced bandwidth throttling settings	76
	Cloud Storage properties	78
	Adding a cloud storage instance	79
	Changing cloud storage host properties	79
	Deleting a cloud storage host instance	81
	About the NetBackup CloudStore Service Container	81
	NetBackup CloudStore Service Container security certificates	82
	NetBackup CloudStore Service Container security modes	83
	NetBackup cloudstore.conf configuration file	84
	Deploying host name-based certificates	87
	Deploying host ID-based certificates	89
	About data compression for cloud backups	90
	About data encryption for cloud storage	91
	About NetBackup KMS for encryption of NetBackup cloud storage	92
	About external KMS for encryption of NetBackup cloud storage	93
	About cloud storage servers	94
	About object size for cloud storage	94
	About the NetBackup media servers for cloud storage	97
	Using media server as NetBackup Cloud primary host	98
	Configuring a storage server for cloud storage	99
	KMS database encryption settings	100
	NetBackup cloud storage server properties	101
	NetBackup cloud storage server bandwidth throttling properties	102
	NetBackup cloud storage server connection properties	106
	NetBackup cloud storage server encryption properties	112
	About cloud storage disk pools	113
	Configuring a disk pool for cloud storage	114
	Saving a record of the KMS key names for NetBackup cloud storage encryption	114
	Adding backup media servers to your cloud environment	117
	Configuring a storage unit for cloud storage	117
	Cloud storage unit properties	118
	Configure a favorable client-to-server ratio	119

	Control backup traffic to the media servers	120
	About NetBackup Accelerator and NetBackup Optimized Synthetic backups	120
	Cloud storage disk pool properties	121
	Certificate validation against Certificate Revocation List (CRL)	123
	Managing Certification Authorities (CA) for NetBackup Cloud	124
Chapter 4	Monitoring and Reporting	127
	Viewing cloud storage job details	127
	Viewing the compression ratio	127
	Displaying KMS key information for cloud storage encryption	128
Chapter 5	Operational notes	131
	NetBackup bpstsinfo command operational notes	131
	Cloud configuration may fail if NetBackup Access Control is enabled	132
	Deleting cloud storage server artifacts	132
	Using <code>csconfig reinitialize</code> to load updated cloud configuration settings	133
	Enabling or disabling communication between primary server and legacy cloud storage media servers	134
Chapter 6	Troubleshooting	136
	About unified logging	136
	About using the <code>vxlogview</code> command to view unified logs	137
	Examples of using <code>vxlogview</code> to view unified logs	139
	About legacy logging	140
	Creating NetBackup log file directories for cloud storage	141
	NetBackup cloud storage log files	142
	Enable libcurl logging	144
	Troubleshooting cloud storage configuration issues	145
	NetBackup Scalable Storage host properties unavailable	145
	Cannot create a cloud storage	146
	Data transfer to cloud storage server fails in the SSL mode	147
	Amazon GovCloud cloud storage configuration fails in non-SSL mode	147
	Data restore from the Google Nearline storage class may fail	147
	Fetching storage regions fails with authentication version V2	148
	Backup from snapshot parent jobs are failing with the status code 160	149

Troubleshooting cloud storage operational issues	149
Cloud storage backups fail	149
Stopping and starting the NetBackup CloudStore Service Container	153
A restart of the nbcssc (on legacy media servers), nbwmc, and nbsl processes reverts all cloudstore.conf settings	154
NetBackup CloudStore Service Container startup and shutdown troubleshooting	154
bptm process takes time to terminate after cancelling GLACIER restore job	155
Handling image cleanup failures for Amazon Glacier vault	155
Cleaning up orphaned archives manually	155
Restoring from Amazon Glacier vault spans more than 24 hours for single fragment	155
Restoring from GLACIER_VAULT takes more than 24 hours for Oracle databases	156
Troubleshooting failures due to missing Amazon IAM permissions	158
Restore job fails if the restore job start time overlaps with the backup job end time	159
Post processing fails for restore from Azure archive	160
Troubleshooting Amazon Snowball and Amazon Snowball Edge issues	160
Index	162

About NetBackup cloud storage

This chapter includes the following topics:

- [About cloud storage features and functionality](#)
- [About the catalog backup of cloud configuration files](#)
- [About support limitations for NetBackup cloud storage](#)

About cloud storage features and functionality

NetBackup Cloud Storage enables you to back up and restore data from cloud Storage as a Service (STaaS) vendors. NetBackup Cloud Storage is integrated with NetBackup OpenStorage.

[Table 1-1](#) outlines the features and functionality NetBackup Cloud Storage delivers.

Table 1-1 Features and functionality

Feature	Details
Cloud storage configuration	NetBackup web UI facilitates the cloud storage setup and storage provisioning.
Compression	NetBackup Cloud Storage Compression compresses the data inline before it is sent to the cloud. The compression feature uses a third-party library called LZO Pro (with compression level 3).

Table 1-1 Features and functionality (*continued*)

Feature	Details
Encryption	NetBackup Cloud Storage Encryption encrypts the data inline before it is sent to the cloud. Encryption interfaces with the NetBackup Key Management Service (KMS) to leverage its ability to manage encryption keys. The encryption feature uses an AES 256 cipher feedback (CFB) mode encryption.
Throttling	NetBackup Cloud Storage throttling controls the data transfer rates between your network and the cloud. The throttling values are set on a per NetBackup media server basis. In certain implementations, you want to limit WAN usage for backups and restores to the cloud. You want to implement this limit so you do not constrain other network activity. Throttling provides a mechanism to the NetBackup administrators to limit NetBackup Cloud Storage traffic. By implementing a limit to cloud WAN traffic, it cannot consume more than the allocated bandwidth. NetBackup Cloud Storage Throttling lets you configure and control the following: ■ Different bandwidth value for both read and write operations. ■ The maximum number of connections that are supported for each cloud provider at any given time. ■ Network bandwidth as a percent of total bandwidth. ■ Network bandwidth per block of time.
Metering	The NetBackup Cloud Storage metering reports enable you to monitor data transfers within NetBackup Cloud Storage. Cloud-based storage is unlike traditional tape or disk media, which use persistent backup images. Your cloud storage vendor calculates cloud-based storage costs per byte stored and per byte transferred. The NetBackup Cloud Storage software uses several techniques to minimize stored and transferred data. With these techniques, traditional catalog-based information about the amount of protected data no longer equates to the amount of data that is stored or transferred. Metering allows installations to monitor the amount of data that is transferred on a per media server basis across one or more cloud-based storage providers.

Table 1-1 Features and functionality (*continued*)

Feature	Details
Cloud Storage service	This is applicable to media server versions 7.7.x to 8.1.2 only. The NetBackup CloudStore Service Container (<code>nbcssc</code>) process performs the following functions: ■ Generates the metering information for the metering plug-in ■ Controls the network bandwidth usage with the help of the throttling plug-in Note: For NetBackup media server versions beyond 8.1.2, these Cloud Storage functions are performed by the NetBackup Service Layer (<code>nbsl</code>) service. On Windows, it is a standard service installed by NetBackup. On UNIX, it runs as a standard daemon. The NetBackup CloudStore Service Container (<code>nbcssc</code>) uses certificate-based authentication. The authentication method used in previous releases (legacy authentication) is disabled by default. It is recommended that you upgrade media servers configured as a cloud storage server to NetBackup 8.1 or later. If you cannot upgrade these servers, use the Enable insecure communication with 8.0 and earlier hosts option on the NetBackup primary server. The option is available in web UI on the Settings > Global Security Settings > Secure Communication tab.
NetBackup Web Management Console	The NetBackup Web Management Console (<code>nbwmc</code>) process manages requests for certificate and host management. This process now also controls the configuration parameters that are related to NetBackup Cloud Storage. The process is installed as a NetBackup service on Windows and runs as a standard daemon on UNIX.

Table 1-1 Features and functionality (*continued*)

Feature	Details
NetBackup Service Layer	The NetBackup Service Layer (<i>nbsl</i>) service facilitates the communication between NetBackup graphical user interface (UI) and the NetBackup logic. This service is also required for Cloud Storage and now performs the following functions: ■ Generates the metering information for the metering plug-in ■ Controls the network bandwidth usage with the help of the throttling plug-in Note: For media server versions 7.7.x to 8.1.2, these Cloud Storage functions are performed by the NetBackup Cloud Storage Service Container (<i>nbcssc</i>).
Storage providers	Cohesity currently supports several cloud storage providers. More information is available about each of these vendors. See “About the cloud storage vendors for NetBackup” on page 14.

About the catalog backup of cloud configuration files

The following cloud configuration files are backed up during the NetBackup catalog backup process:

All `.txt` files in the `meter` directory, which contain intermediate metering data

- `CloudInstance.xml`
- `cloudstore.conf`
- `libstspienencrypt.conf`
- `libstspimetering.conf`
- `libstspithrottling.conf`
- `libstspicloud_provider_name.conf`

All `.conf` files that are specific to the cloud providers that NetBackup supports

The cloud configuration files that are backed up during the catalog backup process reside at the following locations:

Windows `install_path\Veritas\NetBackup\var\global\wmc\cloud`

UNIX `/usr/opensv/var/global/wmc/cloud`

The files `CloudProvider.xml` and `cacert.pem` are at the following location:

Windows `<installed-path>\NetBackup\var\global\cloud`

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

Note: The `cacert.pem` file is not backed up during the NetBackup catalog backup process.

This `cacert.pem` file is a cloud provider-specific file. This file is installed as part of the NetBackup installation. This file includes the well-known public cloud vendor CA certificates used by NetBackup.

About support limitations for NetBackup cloud storage

The following items are some of the limitations of NetBackup cloud storage:

- The cloud vendors do not support optimized duplication.
- The cloud vendors do not support direct to tape (by NDMP).
- The cloud vendors do not support disk volume spanning of backup images.
- If the NetBackup primary server is installed on a platform that NetBackup cloud does not support, you may observe issues in cloud storage server configuration. For the operating systems that NetBackup supports for cloud storage, see the NetBackup operating system compatibility list available through the following URL:
<https://support.cohesity.com/s/article/article-100040093>
- For Hitachi cloud storage, synthetic backups are not successful if you enabled the encryption option. To run the synthetic backups successfully, you need to enable the versioning option for buckets (or namespaces) through the Hitachi cloud portal. For more details on how to enable the versioning option, contact your Hitachi cloud provider.
- Cloud storage servers cannot use the same volume (bucket or container) to store data. You should create a separate volume (bucket or container) for each cloud storage server.
- NetBackup 7.7.1 and later versions support configuring cloud storage using the Frankfurt region.

- In the NetBackup web UI, the following items are displayed only in the English language:
 - All the cloud provider names.
 - Description of the cloud providers.
 - In case of AmazonGov, the following fields: **Certificate File Name**, **Private Key File Name**, **Private Key Passphrase**, **Agency**, **Mission Name**, and **Role**.
 - In case of Openstack Swift, the following fields: **Tenant Type**, **Tenant Value**, **User Type**, **User Domain Type**, **User Domain Value**, **Project Domain Type**, and **Project Domain Value**.
- NetBackup now supports IPv6. The support is available only with all the cloud vendors and proxy server types that support IPv6.

About cloud storage

This chapter includes the following topics:

- [About the cloud storage vendors for NetBackup](#)
- [About the Amazon S3 cloud storage API type](#)
- [About Microsoft Azure cloud storage API type](#)

About the cloud storage vendors for NetBackup

NetBackup supports cloud storage based on the storage API type. All of the cloud vendors that NetBackup supports for cloud storage use one of the supported types. For more information about the storage API types and cloud vendors, see the following:

Cloud storage API types See [Table 2-1](#) on page 15.

The table provides links to the topics that describe the requirements for each storage API type and for the cloud providers who use that storage API type.

Supported cloud vendors Click the following link to identify the list of cloud vendors certified for NetBackup cloud storage and their storage API type: [NetBackup™ Hardware and Cloud Storage Compatibility List \(HCL\)](#)

For configuration help, see the information about their storage API type.

Vendors achieve certification by participating in the Cohesity technology partners program. NetBackup can send backups to the storage that these vendors provide. Cohesity may certify vendors between NetBackup releases. For the vendors that are certified between releases, you must download and install the following configuration and mappings packages:

You can find links to the packages for your release on the NetBackup primary compatibility list landing page:

<https://support.cohesity.com/s/article/article-100040093>

See [Table 2-1](#) on page 15. identifies the cloud storage APIs that are certified for NetBackup cloud storage.

Table 2-1 Supported cloud storage API types for NetBackup

API type	More information
Amazon S3	See “About the Amazon S3 cloud storage API type” on page 15.
Microsoft Azure	See “About Microsoft Azure cloud storage API type” on page 62.

About the Amazon S3 cloud storage API type

NetBackup supports cloud storage from the vendors that use the Amazon S3 storage API for their storage. Information about the requirements and configuration options for the Amazon S3 storage API vendors is provided as follows:

Table 2-2 Amazon S3 storage API type information and topics

Information	Topic
Certified vendors	See “Amazon S3 cloud storage vendors certified for NetBackup” on page 16.
Requirements	See “Amazon S3 storage type requirements” on page 16.
Storage server configuration options	See “Amazon S3 cloud storage provider options” on page 18.
SSL, proxy, and HTTP header options	See “Amazon S3 cloud storage server configuration options” on page 23.
Credential broker options	See “Amazon S3 credentials broker details” on page 26.
Storage classes	See “About Amazon S3 storage classes” on page 28.

Some vendors may support private clouds that use the Amazon S3 storage type API.

See [“About private clouds from Amazon S3-compatible cloud providers”](#) on page 27.

See [“About the cloud storage vendors for NetBackup”](#) on page 14.

Amazon S3 cloud storage vendors certified for NetBackup

Click the following link to identify the vendors who are certified for NetBackup cloud storage: [NetBackup Compatibility List for all Versions](#). Click the link for the **Hardware and Cloud Storage Compatibility List (HCL)** for your version of NetBackup. Inside the HCL, search for **Cloud Storage - Vendor Compatibility** and click **Amazon**.

Vendors achieve certification by participating in the Cohesity Technology Partner Program (VTPP).

Amazon S3 storage type requirements

The following tables describes the details and requirements of Amazon S3 type cloud storage in NetBackup:

Table 2-3 Amazon cloud storage requirements

Requirement	Details
License requirement	You must have a NetBackup license that allows for cloud storage.
Vendor account requirements	You must obtain an account that allows you to create, write to, and read from the storage that your vendor provides.
Buckets	The following are the requirements for the Amazon storage buckets: ■ You can create a maximum of 100 buckets per Amazon account. ■ You can delete empty buckets using the Amazon AWS Management Console. However, you may not be able to reuse the names of the deleted buckets while creating buckets in NetBackup. ■ You can create buckets in any Amazon storage region that NetBackup supports. ■ If the bucket is being used by another user, it is not displayed in the list.

Table 2-3 Amazon cloud storage requirements (continued)

Requirement	Details
Bucket names	It is recommended that you use NetBackup to create the buckets that you use with NetBackup. The Amazon S3 interface may allow the characters that NetBackup does not allow. Consequently, by using NetBackup to create the buckets you can limit the potential problems. The following are the NetBackup requirements for bucket names in the US Standard region. ■ The bucket name must be between 3 and 255 characters. ■ Any of the 26 lowercase (small) letters of the International Standards Organization (ISO) Latin-script alphabet. These are the same lowercase (small) letters as the English alphabet. ■ Any integer from 0 to 9, inclusive. ■ The following character (you cannot use this as the first character in the bucket name): Period (.), underscore (_), and dash (-). Dash - <i>Exception:</i> You cannot use a period (.) if you use SSL for communication. By default, NetBackup uses SSL for communication. See “NetBackup cloud storage server connection properties” on page 106. The buckets are not available for use in NetBackup in the following scenarios: ■ If you have created the buckets in a region that NetBackup does not support. ■ The bucket name does not comply with the bucket naming convention. ■ Given permissions are not sufficient for the bucket. See “Permissions required for Amazon S3 cloud provider user” on page 18.
Number of disk pools	You can create a maximum of 90 disk pools. Attempts to create more than 90 disk pools generate a “failed to create disk volume, invalid request” error message.

Note: You must have SSL enabled to communicate with Amazon AWS. The NetBackup backup job fails with a status code of 87.

See [“About the Amazon S3 cloud storage API type”](#) on page 15.

Permissions required for Amazon S3 cloud provider user

With the Amazon (S3) cloud providers, the following permissions are required to work with NetBackup:

- s3:CreateBucket
- s3:ListAllMyBuckets
- s3:ListBucket
- s3:GetBucketLocation
- s3:GetObject
- s3:PutObject
- s3>DeleteObject
- s3:RestoreObject
- s3:GetBucketObjectLockConfiguration

Amazon S3 cloud storage provider options

[Table 2-4](#) describes the storage server configuration options for Amazon S3.

Table 2-4 Amazon S3 cloud storage provider configuration options

Field name	Required content
Service host	Select the name of the cloud service end point for your vendor from the drop-down list. If the cloud service end point for your vendor does not appear in the drop-down list, you must add a cloud storage instance. See the Add Cloud Storage description in this table.

Table 2-4 Amazon S3 cloud storage provider configuration options
(continued)

Field name	Required content
Storage server name	Displays the default storage server for your vendor. The drop-down list displays only those names that are available for use. If more than one storage server is available, you can select a storage server other than the default one. You can type a different storage server name in the drop-down list, which can be a logical name for the cloud storage. You can create multiple storage servers with different names that refer to the same physical service host for Amazon. If there are no names available in the list, you can create a new storage server name by typing the name in the drop-down list. Note: It is recommended that a storage server name that you add while configuring an Amazon S3-compatible cloud provider should be a logical name and should not match a physical host name. For example: While you add an Amazon GovCloud storage server, avoid using names like 'amazongov.com' or 'amazon123.com'. These servers may be physical hosts, which can cause failures during cloud storage configuration. Instead, use storage server names like 'amazongov1' or 'amazonserver1' and so on. Note: The Add Cloud Storage option is disabled for public clouds. You must use existing cloud storage.
Add Cloud Storage	To configure cloud deployment details, click Storage > Disk storage. The customized cloud deployment refers to the cloud instances that are not already listed in the Service Host drop-down list. After you configure cloud deployment details, the service host appears in the Service Host drop-down list. Once the cloud storage is added, you cannot modify or delete it using web UI. However, you can modify or delete a storage server by using the <code>csconfig</code> command. Note: You can use the NetBackup <code>csconfig -a</code> command to create custom cloud instances for an Amazon S3-compatible cloud provider. You must run the <code>csconfig</code> command before you run the <code>nbdevconfig</code> and <code>tpconfig</code> commands. See the <i>NetBackup Commands Reference Guide</i> for a complete description about these commands. The guide is available through the following URL: https://support.cohesity.com/s/article/article-100040135.html

Table 2-4 Amazon S3 cloud storage provider configuration options
(continued)

Field name	Required content
Media server name	Select a NetBackup media server from the drop-down list. The drop-down list displays only NetBackup 11.2.0.1 and later media servers. In addition, only the media servers that conform to the requirements for cloud storage appear in the drop-down list. The requirements are described in the following topic: See “About the NetBackup media servers for cloud storage” on page 97. The host that you select queries the storage vendor's network for its capabilities and for the available storage. The media server also becomes a data mover for your backups and restores. To support cloud storage, a media server must conform to the following items: ■ The operating system must be supported for cloud storage. For the operating systems that NetBackup supports for cloud storage, see the NetBackup operating system compatibility list available through the following URL: https://support.cohesity.com/s/article/article-100040093 ■ The NetBackup Service Layer (<i>nbsl</i>) service must be running on all the media servers. The NetBackup Web Management Console (<i>nbsmc</i>) must be running on the primary server. ■ For Amazon S3-compatible cloud providers, the media server must run a NetBackup 11.2.0.1 or later release. ■ The NetBackup media servers that you use for cloud storage must be the same NetBackup version as the primary server.
Enter Credentials	<i>Applies to: Amazon GovCloud only.</i> This option is the default selection. Select this option to configure cloud storage server credentials on this wizard panel by entering the access key ID and secret access key.
Use Credentials Broker	<i>Applies to: Amazon GovCloud only.</i> Select this option to configure cloud storage server using credentials broker. If you select this option, you then use the Credentials Broker Details wizard panel that appears next to configure the credentials broker information.

Table 2-4 Amazon S3 cloud storage provider configuration options
(continued)

Field name	Required content
Access key ID	<i>Does not apply for Amazon GovCloud if you select Use Credentials Broker.</i> Enter the access key ID for your vendor account. If you do not have an account, click Create an account with the service provider link.
Secret access key	<i>Does not apply for Amazon GovCloud if you select Use Credentials Broker.</i> Enter the secret access key for your vendor account. It must be 100 or fewer characters.
Use IAM Role(EC2)	NetBackup retrieves the AWS IAM Role name and credentials associated with the EC2 instance. Note: For IAM Role, the selected media server must be hosted on the EC2 instance. See “About using Amazon IAM roles with NetBackup” on page 45.
Advanced Settings	To change SSL, proxy, or HTTP header (server-side encryption or storage class) settings for your cloud storage hosts, click Advanced Settings. See “Amazon S3 cloud storage server configuration options” on page 23.

Amazon S3 cloud storage options

General Settings tab See [Table 2-5](#) on page 22.

Region Settings tab See [Table 2-6](#) on page 23.

Note: If your cloud storage deployment is not configured for multiple regions, you do not need to configure any regions.

Note: To add a cloud storage server in Amazon virtual private cloud (VPC) environment, ensure that you have reviewed the considerations.

See [“Amazon virtual private cloud support with NetBackup”](#) on page 29.

Table 2-5 General settings tab options

Option	Description
Provider type	The cloud storage provider. The following describes the state of this field: ■ Active if you add cloud storage from the Cloud Storage host properties. Select the required provider from the list. ■ Inactive if you add cloud storage from the Storage> Disk storage or change settings from the Cloud Storage host properties. It shows the host that you selected in the wizard or Cloud Storage host properties.
Service host	Enter the cloud service provider host name. If you want to add a public cloud instance, you need to get the service host details from the cloud storage provider. Type the service host details in the text box. If you want to add a cloud storage instance for a private cloud deployment, enter a service host name like 'service.my-cloud.com', in case you can access your cloud provider using the following URL: 'service.my-cloud.com/services/objectstore' For custom instance, to use IPv6 endpoint, you must update or create a new instance with the IPv6 equivalent service host. Note: Do not prefix the service host name with 'http' or 'https'. Note: For VPC in default (US East (N. Virginia)) AWS region, use external-1.amazonaws.com as the service host.
Service endpoint	Enter the cloud service provider endpoint. Service endpoint - Enter the cloud service provider endpoint. For example, '/services/objectstorage' in case your cloud provider service can be accessed using the 'service.my-cloud.com/services/objectstore' URL. You can leave it blank, if the cloud provider service can be accessed directly from the 'service.my-cloud.com' URL.
HTTP port	Enter the HTTP port with which you can access the cloud provider service in a non-secure mode.
HTTPS port	Enter the HTTPS port with which you can access the cloud provider service in a secure mode.

Table 2-5 General settings tab options (*continued*)

Option	Description
Storage server name	Enter a logical name for the cloud storage that you want to configure and access using NetBackup. Note: You can configure multiple storage servers that are associated with the same public or private cloud storage instance.
Endpoint access style	Select the endpoint access style for the cloud service provider. Path style is the default endpoint access style. If your cloud service provider additionally supports virtual hosting of URLs, select Virtual hosted style .

Note: If your cloud storage deployment is not configured for multiple regions, you do not need to configure any regions.

Table 2-6 Region settings tab

Option	Description
Region name	Enter a logical name to identify a specific region where the cloud storage is deployed. For example: East zone.
Location constraint	Enter the location identifier that the cloud provider service uses for any data transfer operations in the associated region. For a public cloud storage, you need to get the location constraint details from the cloud provider. Note: For VPC in default (US East (N. Virginia)) AWS region, use US-east-1 as the location identifier.
Service host	Enter the service host name for the region. The Service endpoint, HTTP port, and HTTPS port information that you have entered in the General settings tab are used while accessing information from any region.
Add	Click Add to add the region.

Amazon S3 cloud storage server configuration options

Note: To access these properties, in the web UI select **Hosts > Host properties**. Select the primary server and click **Edit primary server**. Then click **Cloud Storage**.

The following tables describe the SSL, HTTP header configuration, and proxy server options that are specific to all Amazon S3-compatible cloud providers.

Table 2-7 General settings tab options

Option	Description
Use SSL	Select Use SSL if you want to use the SSL (Secure Sockets Layer) protocol for user authentication or data transfer between NetBackup and cloud storage provider. ■ Authentication only. Select this option if you want to use SSL only at the time of authenticating users while they access the cloud storage. ■ Data transfer. Select this option if you want to use SSL to authenticate users and transfer the data from NetBackup to the cloud storage. Note: NetBackup supports only Certificate Authority (CA) signed certificates while it communicates with cloud storage in the SSL mode. Ensure that the cloud server (public or private) has a CA-signed certificate. If it does not have the CA-signed certificate, data transfer between NetBackup and cloud provider fails in the SSL mode. Note: The FIPS region of Amazon GovCloud cloud provider (that is s3-fips-us-gov-west-1.amazonaws.com) supports only secure 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. Note: Glacier service endpoint for the Amazon GovCloud cloud provider (that is glacier.us-gov-west-1.amazonaws.com) supports only secure mode of communication using the NetBackup GLACIER_VAULT storage class. Therefore, if you disable the Use SSL option while you configure Amazon GovCloud cloud storage with GLACIER_VAULT storage class, the configuration fails.

Table 2-7 General settings tab options (*continued*)

Option	Description
HTTP headers	Specify then appropriate value for the selected HTTP header. Click the Value column to see the drop-down list and select the value. ■ x-amz-server-side-encryption. Select AE256 from the Value drop-down list, if you want to protect data in Amazon S3 cloud storage. AE256 stands for 256-bit Advanced Encryption Standard. By setting the header value to AE256, every object that Amazon S3 cloud storage receives is encrypted before it is stored in the cloud. Amazon S3 server-side encryption uses one of the strongest block ciphers available, that is AE256 to encrypt your data. Additionally, it encrypts the key itself with a primary key that it regularly rotates. Note: If you have already enabled the encryption option when you created the Amazon S3 cloud storage server, you do not need to enable this option. The data is already encrypted before NetBackup sends it over the network. ■ Storage class is configured at the time of creating the storage server. After it is configured, the storage class is non-editable.

Table 2-8 Proxy settings tab options

Option	Description
Use proxy server	Use proxy server option to use proxy server and provide proxy server settings. Once you select the Use Proxy Server option, you can specify the following details: ■ Proxy host. Specify the IP address or name of the proxy server. ■ Proxy port. Specify the 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

Table 2-8 Proxy settings tab options (*continued*)

Option	Description
Use proxy tunneling	You can enable proxy tunneling for HTTP proxy type. After you enable Use proxy tunneling, HTTP CONNECT requests are sent from the cloud media server 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.
Authentication type	You can select one of the following authentication types if you use the HTTP proxy type. ■ None—Authentication is not enabled. A username and password is not required. ■ NTLM—Username and password needed. ■ Basic—Username and password needed. Username. The username of the proxy server. Password. The password can be empty. You can use maximum 256 characters.

See [“About the Amazon S3 cloud storage API type”](#) on page 15.

Amazon S3 credentials broker details

The credentials broker details also appear in Cloud storage server configuration for Amazon GovCloud..

See [“Configuring a storage server for cloud storage”](#) on page 99.

See [“Changing cloud storage host properties”](#) on page 79.

[Table 2-9](#) describes the credential broker options for Amazon GovCloud.

Table 2-9 Credential broker details

Field	Description
Service URL	Enter the service URL. For example: <code>https://hostname:port_number/service_path</code>
Agency	Enter the agency name.
Mission Name	Enter the mission name.

Table 2-9 Credential broker details (*continued*)

Field	Description
Role	Enter the role.
Certificate File Name	Enter the certificate file name.
Private Key File Name	Enter the private key file name.
Private Key Passphrase	Select the check box to specify the private key pass phrase. It must be 100 or fewer characters. The Private Key Passphrase is optional.

Note: The certificate file and the private key file must reside at the following location:

On UNIX - `/usr/opensv/var/global/wmc/cloud`

On Windows - `install_path\Veritas\NetBackup\var\global\wmc\cloud`

Note: For more details on the credentials broker parameters, contact the Cohesity Technical Support team.

See [“About the Amazon S3 cloud storage API type”](#) on page 15.

About private clouds from Amazon S3-compatible cloud providers

NetBackup supports the private clouds or cloud instances from the following Amazon S3-compatible cloud providers:

- Amazon GovCloud
- Cloudian HyperStore
- Hitachi
- Verizon

Before you configure a private cloud in NetBackup, it must be deployed and available.

Use the Advanced Server Configuration dialog box

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

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

The **Create an account with service provider** link on the wizard panel opens a cloud provider webpage in which you can create an account. If you configure a private cloud, that webpage has no value for your configuration process.

About Amazon S3 storage classes

NetBackup supports storage classes for Amazon S3 and Amazon GovCloud. While you configure a cloud storage, you can select a specific storage class that you want to assign to your objects or your data backups. The objects are stored according to their storage classes.

NetBackup supports the following Amazon S3 storage classes:

- **STANDARD**
- **STANDARD_IA** (IA stands for Infrequent Access.)
- **ONEZONE_IA (without Amazon S3 Intelligent Tiering) (IA stands for Infrequent Access.)**
Select the **ONEZONE_IA** (Infrequent Access) storage class to restore less frequently accessed data with single zone resiliency.
- **GLACIER**
Images that are written to Glacier using MSDP direct cloud tiering can be read only by a restore operation. The verify and duplicate operations cannot read the images.
See [“About protecting data in Amazon Glacier”](#) on page 32.
- **GLACIER_VAULT** (Not supported by MSDP direct cloud tiering)
See [“About protecting data in Amazon Glacier vault”](#) on page 36.
- **Glacier Deep Archive**

Images that are written to Glacier Deep Archive using MSDP direct cloud tiering can be read only by a restore operation. The verify and duplicate operations cannot read the images.

See [“About protecting data in Amazon Glacier”](#) on page 32.

- **Amazon S3 Intelligent-Tiering (LIFECYCLE)** (Not supported by MSDP direct cloud tiering)

See [“Protecting data using Amazon S3 Intelligent Tiering \(LIFECYCLE\)”](#) on page 42.

For more about Amazon S3 storage classes, review [Amazon S3 Storage Classes](#).

In the following scenarios, NetBackup assigns the default `STANDARD` storage class to the backups or objects:

- If you do not select a specific storage class while you configure the Amazon S3 cloud storage
- If the backups were configured in an earlier NetBackup version


Note: If you initiate a restore from Glacier or Glacier Deep Archive, NetBackup initiates a warming step. NetBackup does not proceed with the restore until all the data is available in S3 storage to be read.

The warming step is always done if using Amazon. For storage classes other than Glacier and Glacier Deep Archive, the warming step is almost immediate with no meaningful delay. For Glacier and Glacier Deep Archive, the warming step may be immediate if files were previously warmed and are still in S3 Standard storage. However, it may take several minutes, hours, or days depending on settings being used.

Amazon virtual private cloud support with NetBackup

Using NetBackup you can add a new cloud storage in an Amazon virtual private cloud (VPC) environment.

The following diagram illustrates how NetBackup integrates with VPC.



The diagram illustrates the following points:

- You must deploy the media servers within the VPC environment.
- You can deploy the primary server locally or in the VPC environment. Ensure that the primary server is able to communicate with the media servers.
- In the public subnet, PC1 uses both private and elastic IP and has access to the Internet. The media server 1, also has access to the Internet. In a public subnet, you can authenticate and access the storage bucket over Internet or using the VPC endpoint.
- In the private subnet, PC2 uses only private IP and has no access to the Internet. The media server 2, also has no access to the Internet. In a private subnet, you can authenticate and access the storage bucket using the VPC endpoint.
- A VPC is restricted to a specific region.

Considerations for configuring cloud storage server in an Amazon virtual private cloud (VPC) environment

- You need to add a new cloud storage server for the specific region.

- Do not configure multiple regions for one service host.
- When you configure a region for a service host, it must be same as the VPC region; you cannot configure a different region. For example, if you want to add a cloud storage for Singapore region VPC environment, you must configure the service host region to Singapore.
- For VPC in the default (US East (N. Virginia)) AWS region, use **s3-external-1.amazonaws.com** as the service host and **us-east-1** as the location identifier.
- Configure the NetBackup policy to use the media server within the VPC environment.

About protecting data in Amazon for long-term retention

The following Amazon cloud storage options are available for long-term retention of data:

- See [“About protecting data in Amazon Glacier”](#) on page 32.
- See [“About protecting data in Amazon Glacier vault”](#) on page 36.

Difference between GLACIER, GLACIER DEEP ARCHIVE, and GLACIER_VAULT storage classes: When to use what?

Consider the following table when deciding between GLACIER and GLACIER_VAULT storage classes:

GLACIER and GLACIER_DEEP_ARCHIVE storage class	GLACIER_VAULT storage class
GLACIER and GLACIER_DEEP_ARCHIVE storage class corresponds to uploading data through S3 endpoint and transitioning the data to Glacier.	GLACIER_VAULT storage class corresponds to uploading data using Amazon Glacier services to vault.
For GLACIER and GLACIER_DEEP_ARCHIVE storage class, the metadata is stored in STANDARD storage class.	For GLACIER_VAULT storage class, the metadata is stored in STANDARD and GLACIER_VAULT storage classes.
Cost of operation for GLACIER is approximately 2% higher than GLACIER_VAULT.	Cost of operation for GLACIER and GLACIER_VAULT storage class is approximately the same with GLACIER being approximately 2% higher than GLACIER_VAULT.

GLACIER and GLACIER_DEEP_ARCHIVE storage class

Use GLACIER and GLACIER_DEEP_ARCHIVE storage class if you do not plan to use immutable vault lock capability.

GLACIER GLACIER_DEEP_ARCHIVE storage class has a configurable retrieval retention period. Thus, it is useful for restores that may take more time due to size and speed.

As objects get uploaded, Amazon provides visibility for all objects and their storage class property through the Amazon S3 service console. Hence, NetBackup images that are created using GLACIER and GLACIER_DEEP_ARCHIVE storage class have better visibility through the Amazon S3 service console.

There are architectural differences between GLACIER_VAULT storage class (using Amazon Glacier services) and GLACIER and GLACIER_DEEP_ARCHIVE storage class (using Amazon S3 services). This results in difference in speed that must be considered when selecting a storage class.

Storage cleanup handling on failure is better for GLACIER and GLACIER_DEEP_ARCHIVE storage class.

GLACIER_VAULT storage class

Use GLACIER_VAULT storage class if you plan to use the immutable vault lock policy for compliance or to protect your data from ransomware attack.

The retrieval retention period for GLACIER_VAULT storage class is fixed, that is 24 hours. See [“Restoring from Amazon Glacier vault spans more than 24 hours for single fragment”](#) on page 155.

Amazon takes 24 hours to refresh archive inventory. Hence, archives uploaded during backup done using GLACIER_VAULT storage class are reflected in the Amazon Glacier service console only after 24 hours. However, you can get some visibility of backups using the Amazon S3 service console through the metadata generated during the backup. Amazon Glacier service console does not provide any visibility for individual archives.

There are architectural differences between GLACIER_VAULT storage class (using Amazon Glacier services) and GLACIER and GLACIER_DEEP_ARCHIVE storage class (using Amazon S3 services). This results in difference in speed that must be considered when selecting a storage class.

Storage cleanup handling on failure is better for GLACIER storage class as compared to GLACIER_VAULT storage class.

About protecting data in Amazon Glacier

To protect your data for long-term retention you can back up the data to Amazon (AWS) Glacier using NetBackup. Using NetBackup, you can create a storage server with GLACIER or GLACIER_DEEP_ARCHIVE storage class.

To configure a cloud storage server for Amazon GLACIER or DEEP ARCHIVE storage class

- 1 Configure the Amazon GLACIER or GLACIER_DEEP_ARCHIVE cloud storage server.
 See [“Configuring a storage server for cloud storage”](#) on page 99.
- 2 Create a disk pool using the Amazon bucket for GLACIER or GLACIER_DEEP_ARCHIVE storage.
 See [“Configuring a disk pool for cloud storage”](#) on page 114.
- 3 Create a backup policy.
 See the [NetBackup Administrator’s Guide, Volume I](#)

Also ensure that you also have the required permissions. See [“Permissions required for Amazon S3 cloud provider user”](#) on page 18.

To duplicate tape data to Amazon Glacier

Use the `bpduplicate` command to duplicate tape data to Amazon Glacier storage.

Best practices

When you configure a storage server to transition data to Amazon Glacier, consider the following:

- Ensure that GLACIER or GLACIER_DEEP_ARCHIVE is supported for the region to which the bucket belongs.
- For restores, set the retrieval retention period to minimum 3 days.
- Select **True Image Recovery** option wherever possible to reduce time and cost for image imports.

To retrieve the data that is sent to Glacier, there is an inherent time delay of around 4 hours per fragment of the backup image. For phase 2 of image imports, this time delay is prevalent for images in the Glacier storage. However, if you enable **True Image Recovery** in the policy, the time delay for phase 2 imports reduces drastically from 4 hours to a few minutes per fragment. Phase 1 imports are faster, irrespective of whether **True Image Recovery** is enabled or not for the policy.

See the *NetBackup Administrator’s Guide, Volume I* to know more about supported workloads and file systems for **True Image Recovery**.

See the *NetBackup Administrator’s Guide, Volume I* to know more about the phases during image imports.

- You can reduce restore time by parallel restores. For this operation, you use multistreaming to backup which creates multiple images at logical boundaries.

- Workload Granular Recovery (GRT) or VMware Single File Restore (SFR), increases the time-out on the primary, media, and client to more than 5 hours.

Limitations

Consider the following limitations:

- NetBackup Accelerator feature is not supported for policies of the storage units that are created for GLACIER or GLACIER_DEEP_ARCHIVE. Do not select the **Accelerator** check box.

About restoring data from Amazon Glacier

The NetBackup image is stored as set of objects with specified storage class, in this case, GLACIER or GLACIER_DEEP_ARCHIVE storage class. Restore from Amazon Glacier happens in two phases:

- The objects are first retrieved at an internal staging location that Amazon maintains.
- From there, the data is restored at the destination location.

NetBackup supports the following Amazon retrieval types:

- Bulk retrieval, which completes within 5-12 hours.
- Standard retrieval, which completes within 3 – 5 hours.
- Expedited retrieval, which completes within 1-5 minutes.


For more about Amazon S3 storage classes, review [Amazon S3 Storage Classes](#).

Note: If you specify Expedited retrieval, Amazon can sometimes fail the request because of a lack of resources. If this failure happens, you must use Standard retrieval or Bulk retrieval. In this case, the restore job fails (NetBackup status 5: restore failed completely).

The activity monitor displays this message from `bpbrm`: **Image warning failed 503**. The following error is in the `ocsd_storage` log on the MSDP server when MSDP direct cloud tiering is used: **GlacierExpeditedRetrievalNotAvailable: Glacier expedited retrievals are currently not available, please try again later status code: 503**

When you perform a restore, the entire image fragment is restored while only the selected objects are downloaded.

Figure 2-1 Restoring from Amazon Glacier



Note: If you use Glacier with MSDP direct cloud tiering, you can create `GLACIER_RETRIEVAL` touch file on primary server in `/usr/openv/netbackup/bin` directory with one of three strings in it: `bulk`, `standard`, or `expedited`. You can create this touch file if you do not want to use the Bulk retrieval option.

If you use Glacier then you can use `bulk`, `standard`, or `expedited`. If you use `DEEP_ARCHIVE` you can use `bulk` or `standard`. If no string is defined, NetBackup's default is `bulk` if the touch file does not exist.

If you use Glacier with standard, non-deduplication cloud storage servers, only Amazon Standard retrieval is supported.

For more about restoring using Amazon S3, review [Restoring Archived Objects](#).

Considerations with Restore of Image Fragments


Note: This section does not apply to MSDP direct cloud tiering. The section only applies to standard, non-deduplication cloud storage servers.

If the files and folders you want to restore belong to multiple image fragments consider the following:

- One image fragment is retrieved at a time. Only after the selected files and folders part of the first image fragment are downloaded, the next image fragment is retrieved.

- The restore time must be considered depending on the number of image fragments. For example, if the files you want to restore are part of two fragments, an additional 6 - 10 hours are added to the complete restore time.

Figure 2-2 Restoring image fragments for Amazon Glacier



Note: If you cancel a job after the restore retrieval is initiated, cost is incurred for all the objects that are retrieved on the staging location till the point of cancellation.


About protecting data in Amazon Glacier vault

To protect your data for long-term retention with Amazon vault lock policy, you can back up the data to a vault in Amazon Glacier using NetBackup.

When you create a **GLACIER_VAULT** storage class using NetBackup, you specify a vault name and a region in which you want to create the vault.

You can use the Amazon vault lock policy to enforce compliance control on the vault or to make the vault a Write-Once-Read-Many (WORM) device. See the Amazon documentation for more information.

Figure 2-3 Protecting data in Amazon Glacier vault



To configure a cloud storage server for GLACIER_VAULT storage class

- 1 Configure the Amazon GLACIER vault cloud storage server.
See [“Configuring a storage server for cloud storage”](#) on page 99.

Note: Each storage server is associated with only one region.

- 2 Create a disk pool using the Amazon bucket for GLACIER storage.
See [“Configuring a disk pool for cloud storage”](#) on page 114.

Note: If you cannot see the desired vault, it means that either the vault does not have an S3 bucket in the same region as the vault region or the vault does not exist in the region corresponding to the storage server for which you are creating the disk pool.

- 3 Use the Amazon console to create a vault lock policy. See the Amazon documentation for more information.
- 4 Create a backup policy.

Best practices

When you configure a storage server to backup data to a vault in Amazon Glacier, consider the following:

- If you have configured immutable vault lock policy to deny the deletion of archives, Amazon Glacier vault does not allow deletion of archives till the archives are unlocked for deletion. Hence, the retention period configured for a backup policy must be greater than the vault lock period by at least 2 weeks or the maximum time taken to backup or duplicate data to GLACIER_VAULT with retries in your environment. Else, the image cleanup job on image expiration fails. See [“Handling image cleanup failures for Amazon Glacier vault”](#) on page 155.
- It is recommended you use a vault as a secondary target for backing up data.
- If you plan to use the vault lock policy, ensure that you create a vault for each retention level you want to use for the vault.
- Use compression and incremental backups to reduce the size of the data that is stored per backup.
- Select True Image Recovery option wherever possible to reduce time and cost for image imports.

To retrieve the data that is sent to Glacier, there is an inherent time delay of around 4 hours per fragment of the backup image. For phase 2 of image imports, this time delay is prevalent for images in the Glacier storage. However, if you enable **True Image Recovery** in the policy, time spent for phase 2 imports reduces drastically from 4 hours to a few minutes per fragment. Phase 1 imports are faster, irrespective of whether **True Image Recovery** is enabled or not for the policy.

See the *NetBackup Administrator's Guide, Volume I* to know more about supported workloads and file systems for **True Image Recovery**.

See the *NetBackup Administrator's Guide, Volume I* to know more about the phases during image imports.

Limitations

Consider the following limitations:

- NetBackup Accelerator feature is not supported for policies of the storage units that are created for GLACIER_VAULT. Do not select the **Accelerator** check box.
- Glacier endpoint for the Amazon GovCloud cloud provider (that is **glacier.us-gov-west-1.amazonaws.com**) supports only secure mode of communication using the NetBackup **GLACIER_VAULT** storage class. Therefore, if you disable the **Use SSL** option while you configure the Amazon GovCloud cloud storage with **GLACIER_VAULT** storage class, the configuration fails.

Permissions

You must have the following permissions:


- glacier:ListVaults
- glacier:CreateVault
- glacier:DescribeVault
- glacier:UploadArchive
- glacier:DeleteArchive
- glacier:ListJobs
- glacier:DescribeJob
- glacier:InitiateJob
- glacier:GetJobOutput
- Also, ensure that you have the required S3 related IAM USER permissions. See [“Permissions required for Amazon S3 cloud provider user”](#) on page 18.

For permission-related issues, See [“Troubleshooting failures due to missing Amazon IAM permissions”](#) on page 158.

About backing up data to Amazon Glacier vault

When a NetBackup backup job is run to back up data to use the **GLACIER_VAULT** storage class, the data is stored in the vault as a set of archives. The metadata is stored in an S3 bucket as **STANDARD** storage class objects as well as in the Amazon Glacier vault as set of archives.

Figure 2-4 Backing up to Amazon Glacier vault



Considerations:

- If a backup fails due to network issues, the partially backed up data may reside in the vault and hence occupy storage space.

- It is recommended using the Amazon S3 Intelligent Tiering (LIFECYCLE) storage class for moving data from other cloud storage classes to Glacier. See [“Protecting data using Amazon S3 Intelligent Tiering \(LIFECYCLE\)”](#) on page 42. However, to move data from other cloud storage classes to GLACIER_VAULT, you must host a cloud media server and duplicate data through it. This step is done to avoid data-out cost. This workaround also applies to moving data from GLACIER_VAULT storage class to other cloud storage classes.

About restoring data from Amazon Glacier vault


The NetBackup image is stored as a set of data archives in the GLACIER_VAULT storage class. Restore from Amazon Glacier vault happens in two phases.

- The archives are first retrieved at an internal staging location that is maintained by Amazon.
- From there, the data is restored to the destination location.

The staging restore operation takes a minimum of 3 hours to 5 hours. The archives are available at the Amazon staging location for a maximum of 24 hours.

Note: NetBackup supports Amazon Standard retrievals that complete within a minimum of 3 hours to 5 hours. When you perform a restore, the entire image fragment is brought to the staging location while only the selected archives are downloaded.

Figure 2-5 Restoring from Amazon Glacier vault




Considerations with Restore of Image Fragments

If the files and folders you want to restore belong to multiple image fragments, consider the following:

- One image fragment is retrieved at a time. Only after the selected files and folders part of the first image fragment are downloaded, the next image fragment is retrieved.
- The restore time must be considered depending on the number of image fragments. For example, if the files you want to restore are part of two fragments, an additional 6 - 10 hours are added to the complete restore time.

Figure 2-6 Restoring image fragments for Amazon GLACIER_VAULT



Note: If you cancel a job after the restore retrieval is initiated, cost is incurred for all the objects that are retrieved on the staging location till the point of cancellation.

Protecting data using Amazon S3 Intelligent Tiering (LIFECYCLE)

Use the Amazon S3 Intelligent Tiering (LIFECYCLE) storage class to protect your data using cloud tiering. Cloud tiering lets you back up your data to STANDARD or STANDARD_IA storage class and then transition the data to STANDARD_IA or GLACIER storage class. You can configure the storage server properties to determine the number of days the data resides in each storage class. Thus, you can configure your storage server for short-term or long-term data protection.

To configure a cloud storage server for Amazon S3 Intelligent Tiering (LIFECYCLE) storage class

- 1 Configure the Amazon S3 Intelligent Tiering (LIFECYCLE) cloud storage server.
See [“Configuring a storage server for cloud storage”](#) on page 99.
- 2 Configure the storage server properties for the following:

- **AMZ:UPLOAD_CLASS**
- **AMZ:TRANSITION_TO_STANDARD_IA_AFTER**
- **AMZ:TRANSITION_TO_GLACIER_AFTER**

See [“NetBackup cloud storage server connection properties”](#) on page 106.

- 3 Create a disk pool for the Amazon S3 Intelligent Tiering (LIFECYCLE) storage class.

See [“Configuring a disk pool for cloud storage”](#) on page 114.

- 4 Create a backup policy.

Best practices

- Ensure that the selected bucket does not have any existing lifecycle policy.
- If the data is set to transition to GLACIER, consider the following:
 - Ensure that Amazon Glacier is supported for the region to which the bucket belongs.
 - You can use multistreaming to get multiple images at logical boundaries.

Limitations

Consider the following limitations:

- NetBackup Accelerator feature is not supported for policies of the storage units that are created for Amazon S3 Intelligent Tiering (LIFECYCLE). Do not select the **Accelerator** check box.

Permissions

You must have the following permissions:

- Life-cycle policy-related permissions:
 - s3:PutLifecycleConfiguration
 - s3:GetLifecycleConfiguration
- Object tagging permissions
 - s3:PutObjectTagging

Note: The bucket owner has these permissions, by default. The bucket owner can grant these permissions to others by writing an access policy.


- Also ensure that you also have the required IAM USER permissions. See [“Permissions required for Amazon S3 cloud provider user”](#) on page 18.

About backing up data using Amazon S3 Intelligent Tiering (LIFECYCLE) storage class

Initially, the backed up data resides in the storage class determined by the setting **AMZ:UPLOAD_CLASS** in the storage server properties dialog box (default is STANDARD). However, you can configure the duration after which the data transitions to other storage classes by changing the following storage server properties:

- **TRANSITION_TO_STANDARD_IA_AFTER**
- **TRANSITION_TO_GLACIER_AFTER**

Figure 2-7 Back up process for Amazon S3 Intelligent Tiering (LIFECYCLE) storage class with possible configurations



Note: If you want to move data from GLACIER or STANDARD_IA to STANDARD storage class, or GLACIER to STANDARD_IA storage class, you will need to host a cloud media server and duplicate data through it.

After you change the storage server properties and as a new back up job is run per disk pool of the storage server, the new storage server properties get applied to the bucket associated with the disk pool and to the older non-transitioned images in this bucket.

See [“NetBackup cloud storage server connection properties”](#) on page 106.

See [“About protecting data in Amazon Glacier”](#) on page 32.

About restoring data from Amazon S3 Intelligent Tiering (LIFECYCLE) storage class

At the time of restoring, if your data exists in STANDARD or STANDARD_IA storage class, the data is restored to the destination location. However, if the data resides in GLACIER storage class, the data is first retrieved at an internal staging location maintained by Amazon. The data is then restored to the destination location. Hence, the time taken to restore data from STANDARD or STANDARD_IA storage class is much less than the time taken to restore data from GLACIER storage class.

See [“About restoring data from Amazon Glacier”](#) on page 34.

About using Amazon IAM roles with NetBackup

An AWS IAM role is an Amazon Web Services (AWS) identity with the permission policy that determines what tasks an identity is authorized to perform. You can use roles to delegate access to users, applications, or the services that normally don't have access to AWS resources. A role is intended to be assumable by anyone who needs it. If a user assumes a role, temporary security credentials are created dynamically and provided to the user.

For example, an application running on the AWS Elastic Compute Cloud (EC2) instances requires the credentials to access the other AWS services like S3 service. With the traditional approach, you provide the fixed credentials access key and secret access key. With IAM roles, temporary credentials are used to connect to the other AWS services.

Considerations

NetBackup supports the AWS IAM Roles for stream-based backup operations, wherein:

1. NetBackup uses AWS IAM Role that is attached to the AWS EC2 instances on which media server is configured for all S3 storage communications.
2. NetBackup fetches the role name and temporary credentials by connecting to the AWS EC2 metadata.
3. NetBackup primary server can be deployed on AWS EC2 instance or on-premises. You must do the required network settings for communication between the primary and the media servers.
4. TheNetBackup media server that uses the IAM role to backup data to cloud must be deployed on the AWS EC2 instance.
5. AWS IAM Role with required permissions must be attached to the NetBackup media server running on the AWS EC2 instance. See [“Permissions required for Amazon S3 cloud provider user”](#) on page 18.
6. Backup data is stored in S3 storage of the same AWS account where the AWS IAM role is created.
7. NetBackup supports the AWS IAM Role-based authentication for both Amazon and Amazon Gov cloud providers.
8. You can modify existing cloud storage server (alias) to use AWS IAM role for authentication only using the `csconfig` command.
9. Use the AWS Management Console to perform IAM Role allocation, modification, and revocation operations. NetBackup does not store any role-specific information.
10. Ensure that the AWS EC2 instance metadata service (IMDS) is accessible to NetBackup media server. You can verify it using the AWS commands.

You can access instance metadata from a running instance by using IMDSv1 or IMDSv2.

Note: MSDP cloud storage supports IMDSv2. Amazon S3 cloud connector does not support IMDSv2.

IMDSv2 adds defense against open firewalls, reverse proxies, and SSRF vulnerabilities with enhancements to the EC2 instance metadata service. We recommend that you use IMDSv2. For more information see, [AWS documentation](#).

To configure your instance using IMDSv2, see [Configure the instance metadata options](#).

For example,

To get the role name, run:

```
TOKEN=`curl -X PUT "http://169.254.169.254/latest/api/token" -H
"X-aws-ec2-metadata-token-ttl-seconds: 21600"` && curl -H
"X-aws-ec2-metadata-token: $TOKEN" -v
http://169.254.169.254/latest/meta-data/iam/security-credentials/
```


To get the credentials, run:

```
TOKEN=`curl -X PUT "http://169.254.169.254/latest/api/token" -H
"X-aws-ec2-metadata-token-ttl-seconds: 21600"` && curl -H
"X-aws-ec2-metadata-token: $TOKEN" -v
http://169.254.169.254/latest/meta-data/iam/security-credentials/role-name/role-name
```

11. For IPv6 only deployments, AWS IAM Role cannot be used because AWS EC2 instance metadata service is supported only for IPv4.
12. AWS IAM Role is also supported with the MSDP direct cloud tiering storage server.

AWS IAM Role deployment

The following diagram illustrates the deployment:



As the diagram illustrates, to use AWS IAM role with NetBackup:

- NetBackup primary server can be deployed on-premises or in the cloud.
- Backup data is stored in S3 storage of the same AWS account where the AWS IAM role is created.

- AWS IAM role is attached to AWS EC2 instance on which the media server is running.

Note: When role is attached to AWS EC2 instance that has access to S3 storage, NetBackup user doesn't need to provide any credentials.

Tip: You get better performance, if the NetBackup clients are deployed in cloud.

Configuring AWS IAM Role with NetBackup

Using the AWS Management Console and the NetBackup web UI, you can configure AWS IAM Roles with NetBackup.

To configure AWS IAM Role with NetBackup

- 1 Perform the following configurations in the AWS Management Console to use AWS IAM Roles with NetBackup:
 - Create AWS IAM role.
 - Attach role to AWS EC2 instance which will be used as a NetBackup media server.

For guidelines refer to the [technote](#).

- 2 Configure the new cloud storage server to use the AWS IAM role. No credential-specific information is required for using the AWS IAM roles.

See [“Amazon S3 cloud storage provider options”](#) on page 18.

See [“Configuring a storage server for cloud storage”](#) on page 99.

Use the option 'CREDS_ROLE' for credential broker (`-creds_broker`) with the `csconfig` command.

See the [NetBackup Commands Reference Guide](#).

Note: For modifying the existing cloud storage server (alias) to use AWS IAM role for authentication, use only the `csconfig` command.

About NetBackup character restrictions for Amazon S3 cloud connector

NetBackup S3 cloud connector on the S3 compliant cloud storage does not support VMware and Hyper-V backups if the virtual machine display name contains unsupported characters. The unsupported characters are listed in the Object Key Naming guidelines from Amazon S3.

Characters to avoid as per Amazon S3 Object Key Naming guidelines:

The virtual machine display name maps to the key name in Amazon S3 context. Therefore, avoid the following set of characters in a virtual machine display name:

- Backslash \
- Left curly brace {
- Right curly brace }
- Non-printable ASCII characters (128–255 decimal characters)
- Caret ^
- Percent character %
- Grave accent or back tick `
- Right square bracket]
- Left square bracket [
- Quotation marks "
- Tilde ~
- Less Than symbol <
- Greater Than symbol >
- Pound character #
- Vertical bar or pipe |

Characters to avoid as per NetBackup S3 connector guidelines:

Avoid the following set of characters in a virtual machine display name:

- Ampersand &
- Dollar \$
- ASCII character ranges 00–1F hex (0–31 decimal) and 7F (127 decimal)
- At symbol @
- Equals =
- Semicolon ;
- Colon :
- Plus +

- Space (Significant sequences of spaces may be lost in some uses, especially multiple spaces)
- Comma ,
- Question mark ?
- Right round parenthesis)
- Left round parenthesis (

Note: For an updated list of characters to avoid, refer to Amazon S3 documentation.

Protecting data with Amazon Snowball and Amazon Snowball Edge

Amazon Snowball and Amazon Snowball Edge devices can be configured with NetBackup, to backup data to cloud.

The data that is backed up using the Snowball and Snowball Edge devices can be categorized as:

Old data	The backup images that are present in tapes and disks or any other storage media and accumulated over the years.
Live data	Live data The backup data that is generated using daily backups while the Amazon Snowball or the Amazon Snowball Edge device is on-premises. Define storage lifecycle policies for such backups wherein, the actual backup goes to the local storage, and the secondary copy is duplicated to Snowball or Snowball Edge device.

Note: Only STANDARD storage class is supported.

Best practices

Follow these practices when backing up the data to Amazon cloud:

- Plan to keep at least one copy of the data on-premises while data from Snowball or Snowball Edge device is imported to cloud. If backup copy on the Snowball or Snowball device is the only copy you have, use the `bpduplicate` command to make a copy.
 See the [NetBackup Commands Reference Guide](#).
- Verify the imported data in the cloud before discarding (if required) the on-premises backup copy.
- Use the Amazon Snowball and Amazon Snowball Edge device for initial seeding.

- Do not use the buckets for any other purpose before the data is imported to them.
- (For live data) Suspend the duplication operations while the data is in transit and is imported to cloud.
- (For live data) After the data is available in cloud, resume duplication to duplicate the delta data, which was generated on-premises or use another device to transfer it.

Methods

Following are the different methods available for data transfer.

Table 2-10

Device	Methods
Amazon Snowball with NetBackup	Refer to the following topics ■ See “Configuring NetBackup for Amazon Snowball with Amazon Snowball client” on page 52. ■ See “Configuring NetBackup for Amazon Snowball with Amazon S3 API interface” on page 53. ■ See “Configuring SSL for Amazon Snowball and Amazon Snowball Edge” on page 60. ■ After backups are imported into the cloud bucket, you need to perform the post backup procedures. See “Post backup procedures if you have used S3 API interface” on page 61. ■ To improve write performances to the Amazon Snowball device, multiple Amazon S3 adapters can be configured. Also, multiple custom instances can point to the same the device. See “Using multiple Amazon S3 adapters” on page 56.
Amazon Snowball Edge with NetBackup	Refer to the following topics ■ See “Configuring NetBackup with Amazon Snowball Edge with file interface” on page 56. ■ See “Configuring NetBackup for Amazon Snowball Edge with S3 API interface” on page 58. ■ See “Configuring SSL for Amazon Snowball and Amazon Snowball Edge” on page 60. ■ After backups are imported into the cloud bucket, you need to perform the post backup procedures. See “Post backup procedures if you have used S3 API interface” on page 61.

Configuring NetBackup for Amazon Snowball with Amazon Snowball client

In this method data is first staged on the NetBackup media server and then using the Amazon Snowball client, data is moved to the Amazon Snowball device.

Ensure that you have enough space on the file system you plan to use for staging.

To configure NetBackup to transfer data to Amazon Snowball using the Amazon Snowball client

- 1 Create the cloud storage server with default instance.

Note: An Amazon Snowball device can be used to transfer data only from the region from where the device is obtained. Thus, ensure that all the buckets in storage server belong to same region.

Create different bucket(s) for Amazon Snowball when you configure the disk pool. These buckets are used to create an import job in the AWS console.

Note: It is recommended to create the buckets from the NetBackup web UI. However, if you create buckets from the AWS console, ensure that only characters that are supported by NetBackup are used.

See [“Configuring cloud storage in NetBackup”](#) on page 72.

- 2 Create an import job in the AWS console. Select the buckets that were created during the disk pool creation. Refer to the AWS documentation for detailed steps.
- 3 Ensure that the media server has enough space to stage the backup data.
- 4 Update the following storage server properties:
 - `AMZ:OFFLINE_TRANSFER_MODE: FILESYSTEM`
 - `AMZ:TRANSFER_DRIVE_PATH: <absolute path where the data must be backed up>`

Note: Set these properties back to `NONE` after you have transferred the data to the Amazon Snowball device.

See [“NetBackup cloud storage server connection properties”](#) on page 106.

- 5 For live data, create the storage lifecycle policy, backup policy and run the backup for initial seeding.

For old data, use the `bpduplicate` command and duplicate the images on the storage unit.

See the [NetBackup Commands Reference Guide](#).

- 6 Install the Amazon Snowball client on the media server. Refer to the AWS documentation for detailed steps.

Using the Amazon Snowball client, transfer the backup data from the media server to the Amazon Snowball device.

- 7 After the data transfer is complete:
 - Deactivate the backup policy or postpone the secondary operation processing in the SLP till the device is in transit.
 - Set the storage server properties you have configured in step 4 to NONE.
- 8 Ship the device to the cloud vendor. Refer to the AWS documentation for detailed steps.

Example of Amazon client command to move data to Amazon Snowball device

After the backup job is complete, the backup data is staged on the media server. Then run the Amazon Snowball client copy command to transfer the data to the Amazon Snowball device: Following is an example:

```
snowball cp --recursive <TransferDrivePath/MyBucket/Image>  
s3://MyBucket/Logs
```

Refer to the AWS documentation for detailed steps.

Configuring NetBackup for Amazon Snowball with Amazon S3 API interface

When you back up the data to the Amazon Snowball device using the Amazon S3 interface, data is moved directly from the source to the Amazon Snowball device. This process is accomplished using the Amazon S3 APIs.

To configure NetBackup to transfer data to Amazon Snowball using the S3 API interface

- 1 Create a temporary storage server and disk pool to create or list the buckets that you plan to use for the device import job.

Note: It is recommended to create the buckets from the NetBackup web UI. However, if you create buckets from the AWS console, ensure that only characters that NetBackup supports are used.

- 2 Delete the temporary storage server and disk pool.
- 3 Create an import job in the AWS console. Refer to the AWS documentation for detailed steps.
- 4 Install the Amazon Snowball S3 adapter on a different host. Refer to the AWS documentation for detailed steps.
- 5 (Optional) To use SSL protocol for communication with the Amazon Snowball adapter, append the certificate provided to the Amazon Snowball adapter on the command line as it is to `/usr/opensv/var/global/cloud/cacert.pem` file on the media server. Ensure that the format and length of the newly copied certificate matches with the existing certificates in `cacert.pem`.

See [“Configuring SSL for Amazon Snowball and Amazon Snowball Edge”](#) on page 60.

- 6 Add a custom instance for the device.

Set the custom instance’s cloud storage properties with details of the host on which you have installed the Amazon Snowball S3 adapter.

Set the following in the **General Settings** tab:

- Provider type: Amazon or Amazon GovCloud depending upon the endpoint for which you have ordered the device.
- Service host: IP or host name of the adapter
- Service endpoint: Leave blank
- HTTP port: Default is 8080. Or enter the port you have configured.
- HTTPS port: Default is 8443. Or enter the port you have configured.
- Endpoint access style: Path Style

Set the following in the **Region Setting** tab:

- Location constraint: Region from where you have ordered the device.

- Service host: IP or host name of the adapter

Note: An Amazon Snowball Edge device can be used to transfer data only from the region from where the device is obtained. Thus, use the location constraint and service host of the region from where the device is obtained.

See [“Adding a cloud storage instance”](#) on page 79.

- 7 Create a storage server for the device using the custom instance.

See [“Configuring cloud storage in NetBackup”](#) on page 72.

- 8 Update the following storage server property:

```
AMZ:OFFLINE_TRANSFER_MODE: PROVIDER_API
```

See [“NetBackup cloud storage server connection properties”](#) on page 106.

- 9 For live data, create the NetBackup storage lifecycle policy, backup policy and run the backup for initial seeding.

For old data, use the `bpduplicate` command and duplicate the images on the storage unit.

See the [NetBackup Commands Reference Guide](#).

- 10 After the data transfer is complete:

- Deactivate the backup policy or postpone the secondary operation processing in the SLP till the device is in transit.
- Set the storage server properties you have configured to NONE.
- Save the properties. You need this information during the post-backup process.

Take an image capture of storage server properties from Administration console or use `nbdevconfig -getconfig` command. See the [NetBackup Commands Reference Guide](#).

Also, note down the object size that was configured.

- 11 Ship the device to the cloud vendor. Refer to the AWS documentation for detailed steps.

Note: After backups are imported into the cloud bucket, before restore you need to perform the post backup procedures. See [“Post backup procedures if you have used S3 API interface”](#) on page 61.

Using multiple Amazon S3 adapters

To improve write performances to the Amazon Snowball device, multiple Amazon S3 adapters can be configured. Also, multiple custom instances can point to the same the device.

To use multiple Amazon S3 adapter

- 1 For each Amazon Snowball adapter create one custom cloud storage instance.
- 2 Transfer data to the Amazon Snowball device.
- 3 Delete the custom instance with Amazon S3 adapter IP as service host. Run the following command:

```
csconfig cldinstance -r -in <instance-name>
```

See the [NetBackup Command Reference Guide](#).

- 4 Add all the storage servers that are created for the Amazon Snowball device into the default cloud instance (amazon.com). Run the following command:

```
csconfig cldinstance -as -in amazon.com -sts <storage-server-name>
```

- 5 Update the following storage server property:

```
AMZ:OFFLINE_TRANSFER_MODE: NONE
```

See “[NetBackup cloud storage server connection properties](#)” on page 106.

- 6 Change SSL settings (if performed) for the storage servers.

Configuring NetBackup with Amazon Snowball Edge with file interface

When you backup data to the Amazon Snowball Edge device using the file interface, data is moved directly from the source to the Amazon Snowball Edge device.

Recommendation: As a precaution, always create a copy of the backup till the Amazon Snowball Edge device is not imported to cloud.

To configure NetBackup to transfer data to Amazon Snowball Edge using the file interface

- 1 Create the cloud storage server with default instance.

Note: An Amazon Snowball Edge device can be used to transfer data only from the region from where the device is obtained. Thus, ensure that all the buckets in storage server belong to same region.

Create different bucket(s) for Amazon Snowball Edge when you configure the disk pool. These buckets are used to create an import job in the AWS console.

Note: It is recommended to create the buckets from the NetBackup web UI. However, if you create buckets from the AWS console, ensure that only characters that are supported by NetBackup are used.

See [“Configuring cloud storage in NetBackup”](#) on page 72.

- 2 Create an import job in the AWS console. Select the buckets that were created during the disk pool creation. Refer to the AWS documentation for detailed steps.
- 3 Install the Amazon Snowball client on the NetBackup media server.
- 4 Configure the Amazon Snowball Edge device using the Amazon Snowball client.
- 5 Update the following storage server properties:

- `AMZ:OFFLINE_TRANSFER_MODE: FILESYSTEM`
- `AMZ:TRANSFER_DRIVE_PATH: <absolute path of the directory where the file share of the Amazon Snowball Edge device is mounted>`
Mount the root of the file share instead of individual bucket(s) and provide that path to `TRANSFER_DRIVE_PATH`.

Note: Set the property back to `NONE` after you have transferred the data to the Amazon Snowball Edge device.

See [“NetBackup cloud storage server connection properties”](#) on page 106.

- 6 Create the NetBackup storage lifecycle policy and backup policy.
- 7 After the data transfer is complete:

- Deactivate the backup policy or postpone the secondary operation processing in the SLP till the device is in transit.
 - Set the storage server properties you have configured in step 5 to NONE.
 - Rollback the changes you done in step 6.
- 8 Ship the device to the cloud vendor. Refer to the AWS documentation for detailed steps.

Configuring NetBackup for Amazon Snowball Edge with S3 API interface

When you back up the data to the Amazon Snowball Edge device using the S3 interface, data is moved directly from the source to the Amazon Snowball Edge device. This process uses the Amazon S3 adapter.

To configure NetBackup to transfer data to Amazon Snowball Edge using S3 API interface

- 1 Configure the Amazon Snowball Edge device using the Amazon Snowball client.
- 2 Create a temporary storage server and disk pool to create or list the buckets that you plan to use for the device import job.

Note: It is recommended to create the buckets from the NetBackup web UI. However, if you create buckets from the AWS console, ensure that only characters that NetBackup supports are used.

- 3 Delete the temporary storage server and disk pool.
- 4 Create an import job in the AWS console. Refer to the AWS documentation for detailed steps.
- 5 Configure the Amazon Snowball Edge device using the Amazon Snowball client.
- 6 (Optional) To use SSL protocol for communication with the Amazon Snowball Edge, get the certificate using the Amazon snowball client and append the certificate as it is to `/usr/openv/var/global/wmc/cloud/cacert.pem` file on the media server. Ensure that the format and length of the newly copied certificate matches with the existing certificates in `cacert.pem`.

See [“Configuring SSL for Amazon Snowball and Amazon Snowball Edge”](#) on page 60.

- 7 Add a custom instance for the device.

See [“Adding a cloud storage instance”](#) on page 79.

Set the custom instance’s cloud storage properties with details of the host on which you have installed the Amazon Snowball S3 adapter.

Set the following in the **General Settings** tab:

- Provider type: Amazon or Amazon GovCloud depending upon the endpoint for which you have ordered the device.
- Service host: IP or host name
- Service endpoint: Leave blank
- HTTP port: Default is 8080. Or enter the port you have configured.
- HTTPS port: Default is 8443. Or enter the port you have configured.
- Endpoint access style: Path Style

Set the following in the **Region Setting** tab:

- Location constraint: Region from where you have ordered the device.
- Service host: IP or host name

Note: An Amazon Snowball Edge device can be used to transfer data only from the region from where the device is obtained. Thus, use the location constraint and service host of the region from where the device is obtained.

- 8 Create a storage server for the device using the custom instance.

See [“Configuring cloud storage in NetBackup”](#) on page 72.

- 9 Update the following storage server property:

```
AMZ:OFFLINE_TRANSFER_MODE: PROVIDER_API
```

See [“NetBackup cloud storage server connection properties”](#) on page 106.

- 10 For live data, create the storage lifecycle policy, backup policy and run the backup for initial seeding.

For old data, use the `bpduplicate` command and duplicate the images on the storage unit.

See the [NetBackup Commands Reference Guide](#).

- 11 After the data transfer is complete:

- Deactivate the backup policy or postpone the secondary operation processing in the SLP till the device is in transit.
- Set the storage server property you have configured to NONE.

- Save the properties you need this information during the post-backup process.
Take an image capture of storage server properties from web UI or use `nbdevconfig -getconfig` command. See the [NetBackup Commands Reference Guide](#).
Also, note down the object size that was configured.
- 12** Ship the device to the cloud vendor. Refer to the AWS documentation for detailed steps.

Note: After backups are imported into the cloud bucket, before restore you need to perform the post backup procedures. See [“Post backup procedures if you have used S3 API interface”](#) on page 61.

Configuring SSL for Amazon Snowball and Amazon Snowball Edge

To configure SSL for Amazon Snowball

- 1** Ensure that the entries in the `/.aws/snowball/config/snowball-adapter.config` file are correct. Especially, ensure that the host name is set.
- 2** Start the adapter. Following is a sample command:

```
./snowball-adapter -i Snowball IP address -m path to manifest file -u 29 character unlock code --ssl-enabled --aws-secret-key key
```
- 3** Self-signed SSL certificate and key are generated in the `/.aws/snowball/config/` directory.
- 4** Append the certificate provided to the Amazon Snowball adapter on the command line as it is to `/usr/openv/var/global/cloud/cacert.pem` file on the media server. Ensure that the format and length of the newly copied certificate matches with the existing certificates in `cacert.pem`.

To configure SSL for Amazon Snowball Edge

- 1 Lists the certificates available for use. Run the following Amazon Snowball client command:

```
./snowballEdge list-certificates
```
- 2 Obtain the certificate. Run the following Amazon Snowball client command:

```
./snowballEdge get-certificate --certificate-arn arn_value
```
- 3 Append the certificate provided on the command line as it is to `/usr/openv/var/global/cloud/cacert.pem` file on the media server. Ensure that the format and length of the newly copied certificate matches with the existing certificates in `cacert.pem`.

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

Post backup procedures if you have used S3 API interface

After backups are imported into the cloud bucket, perform the following steps before restore:

1. Update the custom instance service host to real endpoint. Also change the HTTP port and region values.
2. (Exception) You cannot update a custom instance for AWS default regions because they are in use by the default-shipped cloud storage instances of NetBackup. Such regions include AWS China Beijing Region, AWS China Ningxia Region, AWS US-East region), AWS GovCloud-US-West and US-East region. For such regions, follow these steps. You can also follow these steps if you encounter an error for unique host name.
 - Keep the saved storage properties handy.
 - Remove the storage server. Run the following command:

```
csconfig cldinstance -rs -in cloud storage instance name -sts storage server name
```


See the [NetBackup Commands Reference Guide](#).
 - Add a new storage server with the same name, under the default storage instance (amazon.com, amazon.cn, amazongov.com, etc.) or the storage instance corresponding to the bucket region. Run the following command to find the instance:

```
csconfig cldinstance -i
```


Run the following command to add the storage server:

```
csconfig cldinstance -as -in Cloud Storage Instance name -sts
storage server name -obj_size size in bytes
```

See the [NetBackup Commands Reference Guide](#).

Ensure that the object size is accurate and same as the storage server that is created.

Also ensure that you have configured the SSL settings as per your requirements.

3. Make sure the SSL setting for storage server as expected. You can verify and update the properties from the **Change Storage Server Connection Properties** dialog box.

See [“To change associated cloud storage server host properties”](#) on page 80.

4. [For Amazon Snowball Edge device only] Update credentials for each storage server with the Amazon account credentials. Run the following command:

```
tpconfig -update -storage_server storage server name -stype
storage server type -sts_user_id [user ID] -password password
```

See the [NetBackup Commands Reference Guide](#).

5. Verify and update the OFFLINE_TRANSFER_MODE storage server property to NONE.
6. Perform the restore and verify the data.
7. Activate policies or activate the secondary operation processing in the SLP .

About Microsoft Azure cloud storage API type

NetBackup supports cloud storage from the vendors that use the Microsoft Azure storage API for their storage. Information about the requirements and configuration options for the Microsoft Azure storage API vendors is provided as follows:

Table 2-11 Microsoft Azure storage API type information and topics

Information	Topic
Certified vendors	See “Microsoft Azure cloud storage vendors certified for NetBackup” on page 63.
Requirements	See “Microsoft Azure storage type requirements” on page 63.
Storage server configuration options	See “Microsoft Azure cloud storage provider options” on page 64.

Table 2-11 Microsoft Azure storage API type information and topics
(continued)

Information	Topic
SSL and proxy options	See "Microsoft Azure advanced server configuration options" on page 66.

Microsoft Azure cloud storage vendors certified for NetBackup

Click the following link to identify the vendors who are certified for NetBackup cloud storage using the Microsoft Azure storage API as of the NetBackup 11.2.0.1 release: [NetBackup™ Enterprise Server and Server 9.0 - 9.x.x Hardware and Cloud Storage Compatibility List \(HCL\)](#).

Vendors achieve certification by participating in the Cohesity Technology Partner Program (VTPP).

Microsoft Azure storage type requirements

[Table 2-12](#) describes the details and requirements of Microsoft Azure cloud storage in NetBackup.

Table 2-12 Microsoft Azure cloud storage requirements

Requirement	Details
License requirement	You must have a NetBackup license that allows for cloud storage.
Microsoft Azure account requirements	You must obtain a Microsoft Azure storage account and at least one storage access key (primary access key or secondary access key).

Table 2-12 Microsoft Azure cloud storage requirements (*continued*)

Requirement	Details
Container names	It is recommended that you use NetBackup to create the container that you use with NetBackup. The following are the NetBackup requirements for container names: ■ Container names must be from 3 through 63 characters long. ■ Container names must start with a letter or number, and can contain only letters, numbers, and the dash (-) character. ■ Every dash (-) character must be immediately preceded and followed by a letter or number; consecutive dashes are not permitted in container names. ■ All letters in a container name must be lowercase. You can refer to the following link: https://msdn.microsoft.com/en-us/library/azure/dd135715.aspx

See “[About Microsoft Azure cloud storage API type](#)” on page 62.

Microsoft Azure cloud storage provider options

[Table 2-13](#) describes the storage server configuration options for Microsoft Azure.

Table 2-13 Microsoft Azure storage server configuration options

Field name	Required content
Service host	Service host is the host name of the cloud service end point of Microsoft Azure. The Service host drop-down list displays part of the service host URL that also comprises Storage Account. Example of a service host URL: <code>storage_account.blob.core.windows.net</code> Note: Based on the region where you have created your storage account - default or China - you should select the service host from the drop-down list.

Table 2-13 Microsoft Azure storage server configuration options (*continued*)

Field name	Required content
Storage server name	Displays the default Azure storage server, which is my-azure. You can select a storage server other than the default one. The drop-down list displays only those names that are available for use. You can type a different storage server name in the drop-down list, which can be a logical name for the cloud storage. You can create multiple storage servers with different names that refer to the same physical service host for Azure. If there are no names available in the list, you can create a new storage server name by typing the name in the drop-down list. Note: It is recommended that a storage server name that you add while configuring an Azure cloud storage should be a logical name and should not match a physical host name. For example: While you add an Azure storage server, avoid using names like 'azure.com' or 'azure123.com'. These servers may be physical hosts, which can cause failures during cloud storage configuration. Instead, use storage server names like 'azure1' or 'azureserver1' and so on.
Media server name	Select a NetBackup media server from the drop-down list. Only the media servers that conform to the requirements for cloud storage appear in the drop-down list. The requirements are described in the following topic: See “About the NetBackup media servers for cloud storage” on page 97. The host that you select queries the storage vendor's network for its capabilities and for the available storage. The media server also becomes a data mover for your backups and restores.
Storage Account	Enter the storage account that you want to use for your cloud backups. For more information about Microsoft Azure storage service, refer to the Microsoft Azure documentation. http://azure.microsoft.com Create the storage account using the following URL: https://portal.azure.com

Table 2-13 Microsoft Azure storage server configuration options (*continued*)

Field name	Required content
Access key	Enter your Azure access key. You can enter the primary access key or the secondary access key. It must be 100 or fewer characters. Refer to the following URL for the access key: https://portal.azure.com
Advanced Settings	To change SSL or proxy settings for Azure, click Advanced Settings . See " Microsoft Azure advanced server configuration options " on page 66.
Configure access tier ACCOUNT_ACCESS_TIER	Select ACCOUNT_ACCESS_TIER option to use the Microsoft Azure account's access tier (Hot or Cool) settings.
Configure access tier ARCHIVE	Select ARCHIVE option for long term retention. See " Protecting data in Microsoft Azure Archive for long-term retention " on page 68.

See "[About Microsoft Azure cloud storage API type](#)" on page 62.

Microsoft Azure advanced server configuration options

The following table describes the SSL and proxy options that are specific to all Microsoft Azure compatible cloud providers. These options appear on the **Advanced Server Configuration** dialog box.

Table 2-14 General settings options

Option	Description
Use SSL	Select this option if you want to use the SSL (Secure Sockets Layer) protocol for user authentication or data transfer between NetBackup and cloud storage provider. ■ Authentication only - Select this option, if you want to use SSL only at the time of authenticating users while they access the cloud storage. ■ Data Transfer - Select this option, if you want to use SSL to authenticate users and transfer the data from NetBackup to the cloud storage. Note: NetBackup supports only Certificate Authority (CA)-signed certificates while it communicates with cloud storage in the SSL mode. Ensure that the cloud server (public or private) has CA-signed certificate. If it does not have the CA-signed certificate, data transfer between NetBackup and cloud provider fails in the SSL mode.

Table 2-15 Proxy Settings tab options

Option	Description
Use Proxy Server	Use Proxy Server option to use proxy server and provide proxy server settings. Once you select the Use Proxy Server option, you can specify the following details: ■ Proxy Host—Specify IP address or name of the proxy server. ■ Proxy Port—Specify port number of the proxy server. ■ Proxy Type— You can select one of the following proxy types: ■ HTTP Note: You need to provide the proxy credentials for HTTP proxy type. ■ SOCKS ■ SOCKS4 ■ SOCKS5 ■ SOCKS4A
Use Proxy Tunneling	You can enable proxy tunneling for HTTP proxy type. After you enable Use Proxy Tunneling, HTTP CONNECT requests are send from the cloud media server to the HTTP proxy server and 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.

Table 2-15 Proxy Settings tab options (*continued*)

Option	Description
Authentication Type	You can select one of the following authentication types if you are using HTTP proxy type. ■ None—Authentication is not enabled. Username and password is not required. ■ NTLM—Username and password needed. ■ Basic—Username and password needed. Username is the username of the proxy server Password can be empty. You can use maximum 256 characters.

See [“About Microsoft Azure cloud storage API type”](#) on page 62.

Protecting data in Microsoft Azure Archive for long-term retention

To protect your data for long-term retention you can back up the data to Microsoft Azure Archive Blob storage using NetBackup. Using NetBackup, you can create a storage server with Archive storage tier.

Note: The Archive storage tier is only available at the blob level and not at the storage account level.

Requirements

Ensure that the following requirements are fulfilled:

- You must have a general-purpose storage V2 to use Azure Archive.

Limitations

Consider the following limitations:

- Accelerator and deduplication are not supported with Azure Archive.
- If restore or cleanup fails, you need to manually set the tier to archive for corresponding blobs.

High-level steps for configurations

1. Configure the Azure Archive cloud storage server.
 See [“Configuring a storage server for cloud storage”](#) on page 99.
2. Create a disk pool with Microsoft Azure Container.
 See [“Configuring a disk pool for cloud storage”](#) on page 114.

3. Storage unit is created using the disk pool.
4. Verify if the `AZR:STORAGE_TIER` property is configured for the storage server.

Note: Once storage server is configured, its `STORAGE_TIER` cannot be changed.

See [“NetBackup cloud storage server properties”](#) on page 101.

5. Use the STU in the backup policy or the service lifecycle policy.


Backing up and restoring data from Azure Archive

During a backup, NetBackup uploads data to the Archive tier.

During a restore, the first image fragments are moved from Archive Tier to HOT tier. Movement of the image fragments takes around 3 hours to 15 hours. After the image fragments are available in the HOT Tier, they are downloaded to local storage. After the restore is complete, the image fragments on the HOT tier are moved back to the Archive Tier.

Note: Image import from Azure Archive storage with TIR is faster.

The following diagram illustrates the restore flow.



Configuring cloud storage in NetBackup

This chapter includes the following topics:

- [Before you begin to configure cloud storage in NetBackup](#)
- [Configuring cloud storage in NetBackup](#)
- [Cloud installation requirements](#)
- [Scalable Storage properties](#)
- [Cloud Storage properties](#)
- [About the NetBackup CloudStore Service Container](#)
- [Deploying host name-based certificates](#)
- [Deploying host ID-based certificates](#)
- [About data compression for cloud backups](#)
- [About data encryption for cloud storage](#)
- [About NetBackup KMS for encryption of NetBackup cloud storage](#)
- [About external KMS for encryption of NetBackup cloud storage](#)
- [About cloud storage servers](#)
- [About object size for cloud storage](#)
- [About the NetBackup media servers for cloud storage](#)
- [Configuring a storage server for cloud storage](#)

- [NetBackup cloud storage server properties](#)
- [About cloud storage disk pools](#)
- [Configuring a disk pool for cloud storage](#)
- [Saving a record of the KMS key names for NetBackup cloud storage encryption](#)
- [Adding backup media servers to your cloud environment](#)
- [Configuring a storage unit for cloud storage](#)
- [About NetBackup Accelerator and NetBackup Optimized Synthetic backups](#)
- [Cloud storage disk pool properties](#)
- [Certificate validation against Certificate Revocation List \(CRL\)](#)
- [Managing Certification Authorities \(CA\) for NetBackup Cloud](#)

Before you begin to configure cloud storage in NetBackup

It is recommended that you do the following before you begin to configure cloud storage in NetBackup:

- Review the NetBackup configuration options for your cloud storage vendor. NetBackup supports cloud storage based on the storage API type, and Cohesity organizes the information that is required to configure cloud storage by API type. The API types, the vendors who use those API types, and links to the required configuration information are in the following topic:
See [“About the cloud storage vendors for NetBackup”](#) on page 14.

Note: Cohesity may certify vendors between NetBackup releases. If your cloud storage vendor is not listed in the NetBackup product documentation, see the following webpage for the most up-to-date list of supported cloud vendors:

<http://www.veritas.com/docs/000115793>

<http://www.veritas.com/docs/000115793>

- Collect the information that is required to configure cloud storage in NetBackup. If you have the required information organized by the NetBackup configuration options, the configuration process may be easier than if you do not.

Configuring cloud storage in NetBackup

This topic describes how to configure cloud storage in NetBackup. [Table 3-1](#) provides an overview of the tasks to configure cloud storage. Follow the steps in the table in sequential order.

The *NetBackup Administrator's Guide, Volume I* describes how to configure a base NetBackup environment. The *NetBackup Administrator's Guide, Volume I* is available through the following URL:

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

Table 3-1 Overview of the NetBackup cloud configuration process

Step	Task	More information
Step 1	Create NetBackup log file directories on the primary server and the media servers	See "NetBackup cloud storage log files" on page 142. See "Creating NetBackup log file directories for cloud storage" on page 141.
Step 2	Review the cloud installation requirements	See "Cloud installation requirements" on page 73.
Step 3	Determine the requirements for provisioning and configuring your cloud storage provider in NetBackup	See "About the cloud storage vendors for NetBackup" on page 14.
Step 4	Configure the global cloud storage host properties as necessary	See "Scalable Storage properties" on page 74.
Step 5	Configure the Cloud Storage properties	Optionally, add a cloud storage service host using the NetBackup host properties. See "Cloud Storage properties" on page 78.
Step 6	Understand the role of the CloudStore Service Container Applicable for media server versions 7.7.x to 8.1.2 only.	See "About the NetBackup CloudStore Service Container" on page 81.
Step 7	Provision a security certificate for authentication on the media servers	See "NetBackup CloudStore Service Container security certificates" on page 82. See "Deploying host name-based certificates" on page 87.

Table 3-1 Overview of the NetBackup cloud configuration process
(continued)

Step	Task	More information
Step 8	Understand key management for encryption	Encryption is optional. See “About data encryption for cloud storage” on page 91. See “About NetBackup KMS for encryption of NetBackup cloud storage” on page 92. See “About external KMS for encryption of NetBackup cloud storage” on page 93.
Step 9	Configure the storage server	See “About cloud storage servers” on page 94. See “Adding a cloud storage instance” on page 79. See “Configuring a storage server for cloud storage” on page 99. See “About object size for cloud storage” on page 94.
Step 10	Configure the disk pool	See “About cloud storage disk pools” on page 113. See “Configuring a disk pool for cloud storage” on page 114.
Step 11	Configure additional storage server properties	See “NetBackup cloud storage server properties” on page 101.
Step 12	Add additional media servers	Adding additional media servers is optional. See “About the NetBackup media servers for cloud storage” on page 97. See “Adding backup media servers to your cloud environment” on page 117.
Step 13	Configure a storage unit	See “Configuring a storage unit for cloud storage” on page 117.

Cloud installation requirements

When you develop a plan to implement a NetBackup Cloud solution, use [Table 3-2](#) to assist with your plan.

Table 3-2 Cloud installation requirements

Requirement	Details
NetBackup media server platform support	For the operating systems that NetBackup supports for cloud storage, see the NetBackup operating system compatibility list available through the following URL: https://support.cohesity.com/s/article/article-100040093 When you install the NetBackup media server software on your host, ensure that you specify the fully-qualified domain name for the NetBackup server name.
Cloud storage provider account	You must have an account created with your preferred cloud storage provider before you configure NetBackup Cloud Storage. Please refer to the list of available NetBackup cloud storage providers. You can create this account in the Cloud Storage Configuration Wizard. See “About the cloud storage vendors for NetBackup” on page 14.
NetBackup cloud storage licensing	NetBackup cloud storage is licensed separately from base NetBackup. The license also enables the Use Accelerator feature on the NetBackup policy Attributes tab. Accelerator increases the speed of full backups for files systems.

Scalable Storage properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the media server. If necessary click **Connect**, then click **Edit media server**. Click **Scalable storage**.

The **Scalable Storage** properties contain information about encryption, metering, bandwidth throttling, and network connections between the NetBackup hosts and your cloud storage provider. These properties appear only if the host is supported for cloud storage. See the *NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List* for your release available through the following URL:

<https://support.cohesity.com/s/article/article-100040093>

The **Scalable storage** properties apply to currently selected media server .

The **Scalable storage** host properties contain the following settings.

Table 3-3 Scalable storage host properties

Property	Description
Key Management Server (KMS) name	If you configured a key management service (KMS) server, the name of the primary server that sends the request to the KMS server is displayed here.
Metering interval	Determines how often NetBackup gathers connection information for reporting purposes. The value is set in seconds. The default setting is 300 seconds (5 minutes). If this value is set to zero, metering is disabled.
Total available bandwidth	Use this value to specify the speed of your connection to the cloud. The value is specified in kilobytes per second. The default value is 102400 KB/sec.
Sampling interval	The time, in seconds, between measurements of bandwidth usage. The larger this value, the less often NetBackup checks to determine the bandwidth in use. If this value is zero, throttling is disabled.
Advanced settings	Expand Advanced settings to configure additional settings for throttling. See “Configuring advanced bandwidth throttling settings” on page 75. See “Advanced bandwidth throttling settings” on page 76.
Maximum concurrent jobs	The default maximum number of concurrent jobs that the media server can run for the cloud storage server. This value applies to the media server, not to the cloud storage server. If you have more than one media server that can connect to the cloud storage server, each media server can have a different value. Therefore, to determine the total number of connections to the cloud storage server, add the values from each media server. If you configure NetBackup to allow more jobs than the number of connections, NetBackup fails any jobs that start after the number of maximum connections is reached. Jobs include both backup and restore jobs. You can configure job limits per backup policy and per storage unit. Note: NetBackup must account for many factors when it starts jobs: the number of concurrent jobs, the number of connections per media server, the number of media servers, and the job load-balancing logic. Therefore, NetBackup may not fail jobs exactly at the maximum number of connections. NetBackup may fail a job when the connection number is slightly less than the maximum, exactly the maximum, or slightly more than the maximum. A value of 100 is generally not needed.

Configuring advanced bandwidth throttling settings

Advanced bandwidth throttling settings let you control various aspects of the connection between the NetBackup hosts and your cloud storage provider.

See [“Scalable Storage properties”](#) on page 74.

To configure advanced bandwidth throttling settings

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the media server.
- 4 If necessary, click **Connect**. Then click **Edit media server**.
- 5 Click **Scalable storage**.
- 6 Expand **Advanced settings**.
- 7 Configure the settings and then click **Save**.

See [“Advanced bandwidth throttling settings”](#) on page 76.

Advanced bandwidth throttling settings

The following table describes the advanced bandwidth throttling settings.

Table 3-4 Advanced throttling configuration settings

Property	Description
Read bandwidth	Use this field to specify the percentage of total bandwidth that read operations can use. Specify a value between 0 and 100. If you enter an incorrect value, an error is generated. If there is insufficient bandwidth to transmit the specified amount of data within a few minutes, restore or replication failures may occur due to time-outs. Consider the total load of simultaneous jobs on multiple media servers when you calculate the required bandwidth. Default value: 100 Possible values: 0 to 100

Table 3-4 Advanced throttling configuration settings (*continued*)

Property	Description
Write bandwidth	Use this field to specify the percentage of total bandwidth that write operations can use. Specify a value between 0 and 100. If you enter an incorrect value, an error is generated. If there is insufficient bandwidth to transmit the specified amount of data within a few minutes, backup failures may occur due to time-outs. Consider the total load of simultaneous jobs on multiple media servers when you calculate the required bandwidth. Default value: 100 Possible values: 0 to 100
Work time	Use this field to specify the time interval that is considered work time for the cloud connection. Specify a start time and end time. Indicate how much bandwidth the cloud connection can use in the Allocated bandwidth field. This value determines how much of the available bandwidth is used for cloud operations in this time window. The value is expressed as a percentage or in kilobytes per second.
Off time	Use this field to specify the time interval that is considered off time for the cloud connection. Specify a start time and end time. Indicate how much bandwidth the cloud connection can use in the Allocated bandwidth field. This value determines how much of the available bandwidth is used for cloud operations in this time window. The value is expressed as a percentage or in kilobytes per second.
Weekend	Specify the start and stop time for the weekend. Indicate how much bandwidth the cloud connection can use in the Allocated bandwidth field. This value determines how much of the available bandwidth is used for cloud operations in this time window. The value is expressed as a percentage or in kilobytes per second.
Read Bandwidth (KB/s)	This field displays how much of the available bandwidth the cloud storage server transmits to a NetBackup media server during each restore job. The value is expressed in kilobytes per second.

Table 3-4 Advanced throttling configuration settings (*continued*)

Property	Description
Write Bandwidth (KB/s)	This field displays how much of the available bandwidth the NetBackup media server transmits to the cloud storage server during backup jobs. The value is expressed in kilobytes per second.

Cloud Storage properties

Note: To access these properties, in the web UI select **Hosts > Host properties**. Select the primary server and click **Edit primary server**. Then click **Cloud Storage**.

The NetBackup **Cloud Storage** properties apply to the currently selected primary server.

The hosts that appear in this **Cloud Storage** list are available to select when you configure a storage server. The **Service provider** type of your cloud vendor determines whether a service host is available or required.

NetBackup includes service hosts for some cloud storage providers. You can add a new host to the **Cloud Storage** list if the **Service provider** type allows it. If you add a host, you also can change its properties or delete it from the **Cloud Storage** list. (You cannot change or delete the information that is included with NetBackup.)

If you do not add a service host to this **Cloud Storage** list, you can add one when you configure the storage server. The **Service provider** type of your cloud vendor determines whether a **Service host name** is available or required.

Cloud Storage host properties contain the following properties:

Table 3-5 Cloud Storage

Property	Description
Cloud Storage	The cloud storage that corresponds to the various cloud service providers that NetBackup supports are listed here. See “Adding a cloud storage instance” on page 79. See “Changing cloud storage host properties” on page 79. See “Deleting a cloud storage host instance” on page 81.
Associated cloud storage servers for <host>	The cloud storage servers that correspond to the selected cloud storage are displayed. See “Changing cloud storage host properties” on page 79.

Adding a cloud storage instance

You may have to add a custom cloud storage instance before you configure a NetBackup cloud storage server. A custom cloud storage allows customization, such as a different service host or other properties.

The cloud storage provider type determines if you have to add a custom cloud storage instance.

See [“About the cloud storage vendors for NetBackup”](#) on page 14.

You can add a custom cloud storage instance as follows:

In the host properties for the primary server	With this method, you add the cloud storage instance before you configure the storage server in NetBackup. Then, the wizard that configures the storage is populated with the instance details. You select the instance when you configure the storage server.
---	--

See [“To add a cloud storage instance in the host properties”](#) on page 79.

Using NetBackup web UI	With this method, you add the instance at the same time as when you configure the storage server in NetBackup. The configured storage is <i>not</i> populated with the instance details until you add.
------------------------	--

See [“Configuring a storage server for cloud storage”](#) on page 99.

To add a cloud storage instance in the host properties

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host Properties**.
- 3 Select the primary server on which to add the cloud storage instance.
- 4 If necessary, click **Connect**. Then click **Edit primary server**.
- 5 Click **Cloud Storage**.
- 6 Click **Add**.
- 7 Configure the settings.

Note: See *NetBackup™ Web UI Administrator's Guide* for details.

- 8 After you configure the settings, click **Save**.

Changing cloud storage host properties

In the **Cloud Storage** properties, you can change the following settings:

Cloud Storage properties	You can change the properties of a host that you add. (You cannot change or delete the properties of the cloud storage providers that are included with NetBackup.) See “To change the Cloud Storage host properties” on page 80.
Associated cloud storage server properties	See “To change associated cloud storage server host properties” on page 80.

How to change cloud storage *server* properties is described in a different topic.

To change the Cloud Storage host properties

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host Properties**.
- 3 Select the primary server on which to add the cloud storage instance.
- 4 If necessary, click **Connect**. Then click **Edit primary server**.
- 5 Click **Cloud Storage**.
- 6 In the **Cloud Storage** list, locate the cloud storage that you want to edit.
- 7 Click **Actions > Edit**.
- 8 Change the properties.
- 9 Click **Save > Save**.

To change associated cloud storage server host properties

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host Properties**.
- 3 Select the primary server on which to add the cloud storage instance.
- 4 If necessary, click **Connect**. Then click **Edit primary server**.
- 5 Click **Cloud Storage**.
- 6 Locate the **Associated cloud storage servers for server** list and then the storage server that you want to edit.
- 7 Click **Edit**.
- 8 Change the properties.

See [“Amazon S3 cloud storage server configuration options”](#) on page 23.

See [“Amazon S3 credentials broker details”](#) on page 26.
- 9 Click **Save > Save**.

Deleting a cloud storage host instance

You can delete your custom cloud storage (cloud instance) in the **Cloud Storage** host properties for the primary server. You cannot delete the cloud storage instances that were delivered with NetBackup.

See [“Cloud Storage properties”](#) on page 78.

To delete a cloud storage host instance

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host Properties**.
- 3 Select the primary server from which to delete the cloud storage.
- 4 If necessary, click **Connect**. Then click **Edit primary server**.
- 5 Click **Cloud Storage**.
- 6 Locate the cloud storage you want to delete.
- 7 Click **Actions > Delete > Delete**.
- 8 Click **Save**.

About the NetBackup CloudStore Service Container

This information is applicable to media server versions 7.7.x to 8.1.2 only.

The NetBackup CloudStore Service Container (`nbcssc`) is a web-based service container that runs on the older media servers that are configured for cloud storage.

This container hosts the throttling service and the metering data collector service.

You can configure the NetBackup CloudStore Service Container behavior in the **Scalable Storage** host properties.

See [“Scalable Storage properties”](#) on page 74.

The port number for the NetBackup CloudStore Service Container service is 5637. Any older media servers that are configured for cloud storage must use this port. Communication with the primary server fails if the older media servers use a different port. Refer to the *NetBackup Network Ports Reference Guide* for more information on the ports that NetBackup uses.

NetBackup uses several methods of security for the NetBackup CloudStore Service Container, as follows:

Security certificates	The NetBackup hosts on which the NetBackup CloudStore Service Container runs must be provisioned with a security certificate or certificates. See “NetBackup CloudStore Service Container security certificates” on page 82. Note: You do not need to generate a security certificate, if you have already generated it before configuring the cloud storage.
Security modes	The NetBackup CloudStore Service Container can run in different security modes. See “NetBackup CloudStore Service Container security modes” on page 83.

See [“About the NetBackup media servers for cloud storage”](#) on page 97.

Note: For NetBackup 8.1.2 and later releases, the `nbcssc` service is no longer deployed. The NetBackup Web Management Console (`nbwmc`) service handles the cloud storage configuration operations and the NetBackup Service Layer (`nbsl`) service handles the throttling service and the metering data collector service functions. For media server versions beyond 8.1.2, authentication is done using host ID-based certificate.

Refer to the *NetBackup Administrator's Guide, Volume I* for more information about these services.

NetBackup CloudStore Service Container security certificates

The NetBackup CloudStore Service Container requires a digital security certificate so that it starts and runs. How the security certificate is provisioned depends on the release level of NetBackup, as follows:

NetBackup 8.2 and later	The NetBackup hosts that run the CloudStore Service Container require a host ID-based certificate. You may have to install the certificate on those hosts. See “Deploying host ID-based certificates” on page 89. If the NetBackup primary server is clustered, you must ensure that the active node and the passive nodes have the host-ID based certificate. See the NetBackup Security and Encryption Guide for more information.
-------------------------	---

NetBackup 8.0 to 8.1.2 The NetBackup hosts that run the CloudStore Service Container require both a host ID-based certificate and a host name-based certificate. You may have to install the certificates on those hosts.

See [“Deploying host name-based certificates”](#) on page 87.

See [“Deploying host ID-based certificates”](#) on page 89.

If the NetBackup primary server is clustered, you must ensure that the active node and the passive nodes have both host named-based and host-ID based certificates. See the [NetBackup Security and Encryption Guide](#) for more information.

Where the media server security certificates reside depend on the release level of NetBackup, as follows:

NetBackup 7.7 to 8.1.2 The certificate name is the host name that you used when you configured the NetBackup media server software on the host. The path for the certificate is as follows, depending on operating system:

- UNIX/Linux: `/usr/openssl/var/vxss/credentials`
- Windows:
`install_dir\Veritas\NetBackup\var\VxSS\credentials`

See [“About the NetBackup CloudStore Service Container”](#) on page 81.

NetBackup CloudStore Service Container security modes

This is applicable only up to NetBackup version 8.1.2.

The NetBackup CloudStore Service Container can run in one of two different modes. The security mode determines how the clients communicate with the service, as follows:

Secure mode	In the default secure mode, the client components must authenticate with the CloudStore Service Container. After authentication, communication occurs over a secure HTTPS channel.
Non-secure mode	The CloudStore Service Container uses non-secure communication. Clients communicate with the server over HTTP with no authentication required.

You can use the `CSSC_IS_SECURE` attribute of the `cloudstore.conf` file to set the security mode. The default value is 64, secure communication.

See [“NetBackup cloudstore.conf configuration file”](#) on page 84.

See [“About the NetBackup CloudStore Service Container”](#) on page 81.

NetBackup cloudstore.conf configuration file

[Table 3-6](#) describes the `cloudstore.conf` configuration file parameters.

The `cloudstore.conf` file is available on the primary server and all the media servers that are installed on the platforms that NetBackup cloud supports.

Note: Before you modify any of the parameters in the `cloudstore.conf` file, you must stop the `nbcssc` service (on media server versions 7.7.x to 8.1.2 only) and the `nbwmc` service (on primary server). Once you modify the parameters, restart these services for the changes to take effect.

The `cloudstore.conf` file resides in the following directories:

- **UNIX:** `/usr/opensv/var/global/cloud`
On media server versions 7.7.x to 8.1.2, the path is:
`/usr/opensv/netbackup/db/cloud`
- **Windows:** `install_path\NetBackup\var\global\cloud`
On media server versions 7.7.x to 8.1.2, the path is:
`install_path\Veritas\NetBackup\db\cloud`

Table 3-6 `cloudstore.conf` configuration file parameters and descriptions

Parameter	Description
CSSC_VERSION	It is not recommended to modify this value. Specifies the version of <code>cloudstore.conf</code> file. The default value is 2.
CSSC_PLUGIN_PATH	It is not recommended to modify this value. Specifies the path where NetBackup cloud storage plug-ins are installed. The default path is as follows: On Windows: <code>install_path\Veritas\NetBackup\bin\ost-plugins</code> On UNIX: <code>/usr/opensv/lib/ost-plugins</code>

Table 3-6 `cloudstore.conf` configuration file parameters and descriptions
(continued)

Parameter	Description
CSSC_PORT	This setting is applicable to media server versions 7.7.x to 8.1.2 only. Specifies the port number for the CloudStore Service Container (<code>nbcssc</code>). Specify the value as 5637. This port is used to provide back-level media server support for the older media servers that are configured for cloud storage. Ensure that the older media servers use this port. Communication with the primary server fails if the older media servers use a different port.
CSSC_LOG_DIR	Specifies the directory path where <code>csconfig</code>, <code>nbcldutil</code>, and cloud plug-ins generate log files. The default path is as follows: On Windows: <code>install_path\Veritas\NetBackup\logs\nbcssc</code> On UNIX: <code>/usr/opensv/netbackup/logs/nbcssc</code> Note: For media server versions 7.7.x to 8.1.2, the <code>nbcssc</code> service uses this path for log files.
CSSC_LOG_FILE	This setting is applicable only up to NetBackup release 8.1.2. Specifies the file name that the <code>nbcssc</code> service uses to write its logs. The default value is empty, which means that the NetBackup logging mechanism determines the log file name.
CSCONFIG_LOG_FILE	Specifies the file name that the <code>csconfig</code> utility uses to write its logs. The default value is empty, which means that the NetBackup logging mechanism determines the log file name.
CSSC_IS_SECURE	Specifies if the <code>nbcssc</code> service runs in secure (value 64) or non-secure mode (value 0). The default value is 64.

Table 3-6 `cloudstore.conf` configuration file parameters and descriptions
(continued)

Parameter	Description
<code>CSSC_CIPHER_LIST</code>	Specifies the cipher list that NetBackup uses for the following purposes: ■ The cloud primary host's cipher is used for communication with the cloud service provider. ■ The media server cipher is used for communicating with the cloud primary host's <code>nbwmc</code> service and with the cloud service provider. It is recommended that you do not modify this value. Depending on the purpose for customizing the cipher list, you must modify the cipher list in the <code>cloudstore.conf</code> on the primary server and the media servers. Note: If the cipher list is invalid, the customized cipher list is replaced by the default cipher list. The default value is <code>AES:!aNULL:@STRENGTH</code>.
<code>CSSC_LOG_LEVEL</code>	Specifies the log level for <code>csconfig</code> and <code>nbclutil</code> CLI utility logging. Value 0 indicates that the logging is disabled and a non-zero value indicates that the logging is enabled. The default value is 0.
<code>CSSC_MASTER_PORT</code>	This setting is applicable for media server versions 7.7.x to 8.1.2 only. It is not applicable for NetBackup primary and media server versions 8.2 and later. This parameter value must be set to 5637. This port is used to provide back-level support for the older media servers that are configured for cloud storage. Ensure that the older media servers use this port. Communication with the primary server fails if the older media servers use a different port.
<code>CSSC_MASTER_NAME</code>	Specifies the NetBackup primary server name. This entry indicates that the <code>nbwmc</code> service runs on this host. It processes all cloud provider-specific requests based on the <code>CloudProvider.xml</code> and <code>CloudInstance.xml</code> files.

Table 3-6 `cloudstore.conf` configuration file parameters and descriptions
(continued)

Parameter	Description
CSSC_ALLOW_LEGACY_AUTH	Specifies if the primary server can communicate with the legacy media servers that are configured for cloud storage. Only media server versions 7.7.x to 8.1.2 are supported. The value 1 (default value) indicates that the communication is enabled while the value 0 means that the communication is disabled. Use this parameter with the Enable insecure communication with 8.0 and earlier hosts option available in the NetBackup web UI (Settings > Global security > Secure communication). The GUI option lets you enable or disable primary server communication with all back-level legacy media servers. It works as an all or none kind of a setting and is not specific to cloud storage media servers. This parameter provides that additional level of control for the cloud. You can use this setting to enable or disable primary server communication with back-level cloud storage media servers explicitly. For example, if the GUI option is enabled (default value) and this parameter value is set to 0, the NetBackup primary server continues to work with supported back-level media servers as other storage servers. However, any legacy cloud storage media servers that use the older method of communication using hard-coded credentials are blocked altogether, thus increasing the security of your NetBackup environment. Note: This parameter value has no effect if the GUI option is disabled. If you modify this parameter value, you must restart the NetBackup Web Management Console (nbwmc) service for the changes to take effect.

Deploying host name-based certificates

This is applicable for media server versions 7.7.x to 8.1.2 only.

You can deploy the required host name-based security certificate for the NetBackup media servers that you use for cloud storage. Each media server that you use for cloud storage runs the NetBackup CloudStore Service Container.

See [“About the NetBackup CloudStore Service Container”](#) on page 81.

You can deploy a certificate for an individual media server or for all media servers. Media servers that you use for cloud storage must have a host name-based security certificate.

Note: Deploying a host name-based certificate is a one-time activity for a host. If a host name-based certificate was deployed for an earlier release or for a hotfix, it does not need to be done again.

Ensure the following before you deploy a host-name based certificate:

- All nodes of the cluster have a host ID-based certificate.
- All Fully Qualified Domain Names (FQHN) and short names for the cluster nodes are mapped to their respective host IDs.

Deploying a host name-based certificate on media servers

This procedure works well when you deploy host name-based security certificates to many hosts at one time. As with NetBackup deployment in general, this method assumes that the network is secure.

To deploy a host name-based security certificate for media servers

- 1 Run the following command on the primary server, depending on your environment. Specify the name of an individual media server or specify `-AllMediaServers`.

On Windows: `install_path\NetBackup\bin\admincmd\bpnbaz
-ProvisionCert host_name|-AllMediaServers`

On UNIX: `/usr/openv/netbackup/bin/admincmd/bpnbaz -ProvisionCert
host_name|-AllMediaServers`

NetBackup appliance (as a NetBackupCLI user): `bpnbaz -ProvisionCert
Media_server_name`

- 2 Restart the NetBackup Service Layer (nbsl) service on the media server.

Note: In you use dynamic IPs on the hosts (DHCP), ensure that the host name and the IP address are correctly listed on the primary server. To do so, run the following NetBackup `bpclient` command on the primary server:

On Windows: `Install_path\NetBackup\bin\admincmd\bpclient -L -All`

On UNIX: `/usr/openv/netbackup/bin/admincmd/bpclient -L -All`

Deploying host ID-based certificates

Depending on the certificate deployment security level, a non-primary host may require an authorization token before it can obtain a host ID-based certificate from the Certificate Authority (primary server). When certificates are not deployed automatically, they must be deployed manually by the administrator on a NetBackup host using the `nbcertcmd` command.

The following topic describes the deployment levels and whether the level requires an authorization token.

Deploying when no token is needed

Use the following procedure when the security level is such that a host administrator can deploy a certificate on a non-primary host without requiring an authorization token.

To generate and deploy a host ID-based certificate when no token is needed

- 1 The host administrator runs the following command on the non-primary host to establish that the primary server can be trusted:

```
nbcertcmd -getCACertificate
```

- 2 Run the following command on the non-primary host:

```
nbcertcmd -getCertificate
```

Note: To communicate with multiple NetBackup domains, the administrator of the host must request a certificate from each primary server using the `-server` option.

Run the following command to get a certificate from a specific primary server:

```
nbcertcmd -getCertificate -server primary_server_name
```

- 3 To verify that the certificate is deployed on the host, run the following command:

```
nbcertcmd -listCertDetails
```

Deploying when a token is needed

Use the following procedure when the security level is such that a host requires an authorization token before it can deploy a host ID-based certificate from the CA.

To generate and deploy a host ID-based certificate when a token is required

- 1 The host administrator must have obtained the authorization token value from the CA before proceeding. The token may be conveyed to the administrator by email, by file, or verbally, depending on the various security guidelines of the environment.

- 2 Run the following command on the non-primary host to establish that the primary server can be trusted:

```
nbcertcmd -getCACertificate
```

- 3 Run the following command on the non-primary host and enter the token when prompted:

```
nbcertcmd -getCertificate -token
```

Note: To communicate with multiple NetBackup domains, the administrator of the host must request a certificate from each primary server using the `-server` option.

If the administrator obtained the token in a file, enter the following:

```
nbcertcmd -getCertificate -file authorization_token_file
```

- 4 To verify that the certificate is deployed on the host, run the following command:

```
nbcertcmd -listCertDetails
```

Use the `-cluster` option to display cluster certificates.

About data compression for cloud backups

In NetBackup, you can compress your data before you send it to cloud storage server.

You can enable data compression on the NetBackup media server while you configure your cloud storage server using NetBackup web UI.

See [“Configuring a storage server for cloud storage”](#) on page 99.

Note: After you have enabled the data compression during the cloud storage configuration, you cannot disable it.

Important notes about data compression in NetBackup

- NetBackup uses a third-party library, LZO Pro, with compression level 3. The `bptm` logs provide information of the compression ratio of your data after the backup is taken in the cloud storage.
 See [“Viewing the compression ratio”](#) on page 127.
- NetBackup compresses the data in chunks of 256 KB.
- NetBackup Accelerator and True Image Restore (TIR) with move detection is supported with compression.
- The backup data is compressed before it is transmitted to the cloud storage server. If both the compression and the encryption options are selected, the data is compressed before it is encrypted.
- Data compression reduces the backup time and the data size based on how much the data is compressible. Although you may notice reduced bandwidth utilization when you compare it with the data without compression.
- Performance of the data compression is reduced, if the data is incompressible. Therefore, it is not recommended to enable compression for backing up incompressible data such as policy data and so on.
- It is not recommended to use the same bucket with storage servers of different types.
- You must not use client-side compression along with storage server-side compression.
- You cannot change the compression configuration settings (enable/disable) after the storage server is created.

About data encryption for cloud storage

You can encrypt your data before you send it to the cloud. You can configure encryption while configuring the storage.

NetBackup uses NetBackup Key Management Service (NetBackup KMS) and external key management service (external KMS) for managing data encryption in case of cloud disk storage.

See [“About NetBackup KMS for encryption of NetBackup cloud storage”](#) on page 92.

See [“About external KMS for encryption of NetBackup cloud storage”](#) on page 93.

More information about NetBackup KMS and external KMS is available.

See the [NetBackup Security and Encryption Guide](#) for more information.:

About NetBackup KMS for encryption of NetBackup cloud storage

NetBackup uses NetBackup Key Management Service (NetBackup KMS) to manage the keys for the data encryption for disk storage. NetBackup KMS is a NetBackup primary server-based symmetric key management service. The service runs on the NetBackup primary server. An additional license is not required to use the NetBackup KMS functionality. NetBackup uses NetBackup KMS to manage the encryption keys for cloud storage.

See [“About data encryption for cloud storage”](#) on page 91.

You need to provide KMS and key-specific information when you enable the **Encryption through Cloud Storage Server Configuration Wizard** and configure disk pool using the **Disk Pool Configuration Wizard**. Key-specific information is based on the KMS server configuration. If a KMS server is not configured, NetBackup KMS is by default configured as a KMS server as part of the encryption setting for the cloud storage server.

The following table describes the keys that are required for the NetBackup KMS database. You can enter the pass phrases for these keys when you use the **Cloud Storage Server Configuration Wizard**.

Table 3-7 Encryption keys required for the KMS database

Key	Description
Host Master Key	The Host Master Key protects the key database. The Host Master Key requires a pass phrase and an ID. NetBackup KMS uses the pass phrase to generate the key.
Key Protection Key	A Key Protection Key protects individual records in the key database. The Key Protection Key requires a pass phrase and an ID. NetBackup KMS uses the pass phrase to generate the key.

The following table describes the encryption keys that are required for each storage server and volume combination. If you specify encryption when you configured the cloud storage server, you must configure a pass phrases for the key group for the storage volumes. You enter the pass phrase for these keys when you use the **Disk Pool Configuration Wizard**.

Table 3-8 Encryption keys and key records for each storage server and volume combination

Item	Description
Key group key	A key group key protects the key group. Each storage server and volume combination requires a key group, and each key group key requires a pass phrase. The key group name must use the format for the storage type that is described as follows: For cloud storage, the following is the format: <pre>storage_server_name:volume_name</pre> The following items describe the requirements for the key group name components for cloud storage: ■ <i>storage_server_name</i> : You must use the same name that you use for the storage server. The name can be a fully-qualified domain name or a short name, but it must be the same as the storage server. ■ The colon (:) is required after the <i>storage_server_name</i>. ■ <i>volume_name</i> : You must specify the LSU name that the storage vendor exposes to NetBackup. The Disk Pool Configuration Wizard conforms to this format when it creates a key group.
Key record	Each key group that you create requires a key record. A key record stores the actual key that protects the data for the storage server and volume. A name for the key record is optional. If you use a key name, you can use any name. It is recommended that you use the same name as the volume name. The Disk Pool Configuration Wizard does not prompt for a key record key; it uses the volume name as the key name.

More information about NetBackup KMS and external KMS is available in the [NetBackup Security and Encryption Guide](#).

About external KMS for encryption of NetBackup cloud storage

NetBackup supports keys from external key management service (external KMS) server in case of cloud storage.

Symmetric encryption key is required for each storage server and volume combination. Symmetric encryption key is not created on the external KMS server for each storage server and volume combination. You need to ensure that a

Symmetric encryption key already exists on the external KMS server with a custom attribute with value of key group in the 'storage_server_name:volume_name' format.

More information about external KMS is available in the [NetBackup Security and Encryption Guide](#).

About cloud storage servers

A storage server is an entity that writes data to and reads data from the storage. In case of cloud storage server, it is a host or an end point that cloud vendor exposes to perform backup operations using NetBackup media server(s). You can use any logical name to identify the cloud storage when you configure cloud storage server in NetBackup.

When you configure a cloud storage server, it inherits the NetBackup Scalable Storage properties.

See [“Scalable Storage properties”](#) on page 74.

After you configure the storage server, you can change the properties of the storage server.

NetBackup media servers back up the clients and send the data to the storage server.

See [“About the NetBackup media servers for cloud storage”](#) on page 97.

About object size for cloud storage

During backup, NetBackup divides the backup image data into chunks called objects. PUT request is made for each object to move it to the cloud storage.

By setting a custom Object Size, you can control the amount of PUT and GET requests that are sent to and from the cloud storage. The reduced number of PUT and GET requests help in reducing the total charges that are incurred for the requests.

During the creation of a cloud storage server, you can specify a custom value for the Object Size. Consider the cloud storage provider, hardware, infrastructure, expected performance, and other factors for deciding the value. Once you set the Object Size for a cloud storage server, you cannot change the value. If you want to set a different Object Size, you must recreate the cloud storage server.

See [“Configuring a storage server for cloud storage”](#) on page 99.

Guidelines for selecting the Object Size


The combination of object size, number of parallel connections, and the read or write buffer size contribute to the performance of NetBackup in the cloud.

To enhance the performance of backup and restore operations, NetBackup uses multiple parallel connections into cloud storage. The performance of NetBackup depends on the number of parallel connections. Number of parallel connections are derived from the read or write buffer size and the object size.

Read or Write buffer size (user set) ÷ Object Size (user set) = Number of parallel connections (derived). The following diagram illustrates how these factors are related:

The following diagram illustrates how these factors are related:

Figure 3-1 Object size



- Consider the following factors when deciding the number of parallel connections:
 - The maximum number of parallel connections the cloud storage provider permits.
 - Network bandwidth availability between NetBackup and the cloud storage environment.
 - System memory availability on the NetBackup host.
- If you increase the object size, the number of parallel connections reduce. The number of parallel connections affect the upload and the download rate.
- If you increase the read or write buffer size, the number of parallel connections increase. Similarly, if you want lesser number of parallel connections, you can

reduce the read or write buffer size. However, you must consider the network bandwidth and the system memory availability.

- Cloud providers charge for the number of PUT and GET requests that are initiated during a backup or restore process. The smaller the object size, higher the number of PUT or GET requests, and therefore, higher charges are incurred.
- In case of temporary failures with data transfer, NetBackup performs multiple retries for transferring the failed objects. If the failures persist, the complete object is transferred again. Also, with higher latency and higher packet loss, the performance might reduce. To handle the latency and the packet loss issues, increasing the number of parallel connections can be helpful.
- NetBackup has some time-outs on the client side. If the upload operation takes more time (due to big object size) than the minimum derived NetBackup data transfer rate, there can be failures with NetBackup.
- For legacy environments without deduplication support, if the number of connections are less, parallel downloads are less compared to older number of connections.

For example, while restoring from back-level images (8.0 and earlier), where the object size is 1MB, the buffer of 16 MB (for one connection) is not completely used while also consuming memory. With the increased object size, there is a restriction on number of connections due to the available read or write buffer size memory.

Current default settings

The default settings are as follows:

Table 3-9 Current default settings

Cloud storage provider	Object size	Default read or write buffer size
Amazon S3/Amazon GovCloud	16 MB (fixed)	400 MB (configurable between 16 MB to 1 GB)
Azure	4 MB (fixed)	400 MB (configurable between 4 MB to 1 GB)

About the NetBackup media servers for cloud storage

The NetBackup media servers that you use for cloud storage backup the NetBackup clients and then send that backup data to the cloud storage server. The storage server then writes the data to storage.

See [“About cloud storage servers”](#) on page 94.

The NetBackup media servers also can move data back to primary storage (the client) during restores and from secondary storage to tertiary storage during duplication. These media servers are also known as *data movers*. They host a software plug in that they use to communicate with the storage implementation.

When you configure a cloud storage server, the media server that you specify in the wizard or on the command line becomes a cloud storage data mover.

See [“Configuring a storage server for cloud storage”](#) on page 99.

You can add additional media servers to backup clients. They can help balance the load of the backups that you send to the cloud storage.

See [“Adding backup media servers to your cloud environment”](#) on page 117.

You can control which data movers are used for backups and duplications when you configure NetBackup storage units.

See [“Configuring a storage unit for cloud storage”](#) on page 117.

You can configure a cloud media server as a cloud primary host.

See [“Using media server as NetBackup Cloud primary host”](#) on page 98.

To support cloud storage, a media server must conform to the following items:

- The operating system must be supported for cloud storage. For the operating systems that NetBackup supports for cloud storage, see the NetBackup operating system compatibility list available through the following URL:
<https://support.cohesity.com/s/article/article-100040093>
<https://support.cohesity.com/s/article/article-100040093>
- On media server versions 7.7.x to 8.1.2, the NetBackup Cloud Storage Service Container (`nbcssc`) must be running.
 See [“About the NetBackup CloudStore Service Container”](#) on page 81.
- The NetBackup media servers that you use for cloud storage must be the same NetBackup version as the primary server.

Using media server as NetBackup Cloud primary host

These steps are applicable to media server versions up to 8.1.2 only.

You must perform this procedure for all the operating systems those are not supported by NetBackup cloud.

See the NetBackup hardware compatibility list for your release available through the following URL:

<https://support.cohesity.com/s/article/article-100040093>

For disaster recovery, you must take a manual backup of the following files from the media server that you have configured as NetBackup cloud primary host:

- `CloudProvider.xml`
- `CloudInstance.xml`

To use media server as NetBackup cloud primary host

- 1 Identify one of the NetBackup cloud media servers as a cloud primary host.

Choose a media server that has same NetBackup primary server version. Do not use a media server with different version.

Note: The media server does not hold the primary copy of the `CloudProvider.xml` file which all the media servers require while configuring the cloud storage and for running operations such as backup, restore, and so on.

- 2 Run the following commands on all the NetBackup cloud media servers including the one that is selected as the cloud primary host:

```
nbcssc -t -a Netbackup
```

```
nbcssc -s -a Netbackup -m cloud_master_host -f
```

For information on the command, see [NetBackup Commands Reference Guide](#).

- 3 Ensure that the values of **CSSC_PORT** and **CSSC_IS_SECURE** as mentioned in `cloudstore.conf` file from cloud primary host are copied as **CSSC_MASTER_PORT** and **CSSC_MASTER_IS_SECURE** in `cloudstore.conf` file on all other NetBackup cloud media servers.

After you select a cloud primary host, do not change the name again to point to another media server. If you need to do so, contact Cohesity Technical Support.

Additional task post disaster recovery

For a cloud storage server that uses proxy server, you must update the proxy credentials.

- To perform the task using the NetBackup Administrators Console, see [“Changing cloud storage host properties”](#) on page 79.
- To perform the task using the commands, run the following:

```
csconfig cldinstance -us -in instance_name -sts storage_server_name  
-pxtype proxy_type -pxhost proxy_host -pxport proxy_port  
-pxautth_type proxy_auth_type -pxtunnel proxytunnel_usage
```

For information on the command, see [NetBackup Commands Reference Guide](#).

Additional task post primary server upgrade

This is applicable for a NetBackup environment where a primary server is running on an unsupported operating system such as Solaris x86 or Windows Server 2008, and the media server is promoted as a Cloud Primary host.

After upgrading the primary server, if you plan to perform a rolling upgrade on the media server, then there are some additional post-upgrade steps that must be performed to ensure that the cloud storage server works seamlessly after the media server upgrade is completed.

Refer to the following technote for more details:

<https://support.cohesity.com/s/article/article-100044766>

Configuring a storage server for cloud storage

Configure in this context means to configure a host as a storage server that can write to and read from the cloud storage.

See [“About cloud storage servers”](#) on page 94.

NetBackup UI also lets you enable encryption and configure corresponding parameters for the NetBackup Key Management Service (NetBackup KMS) server if no KMS server is configured.

See [“About data encryption for cloud storage”](#) on page 91.

If data encryption and NetBackup KMS are configured, it is recommended that you save a record of key names.

See [“Saving a record of the KMS key names for NetBackup cloud storage encryption”](#) on page 114.

If you configure a storage server by using CLI, you must run `csconfig` command before running `nbdevconfig` and `tpconfig` commands.

See the [NetBackup Commands Reference Guide](#).

The NetBackup media server that you select during the configuration process must conform to the requirements for cloud storage.

See “[About the NetBackup media servers for cloud storage](#)” on page 97.

Refer to the *NetBackup Web UI Administrators Guide* for details of configuring a cloud storage server.

KMS database encryption settings

This section describes the settings to configure the NetBackup Key Management Service database and the encryption keys for your cloud storage. This information protects the database that contains the keys that NetBackup uses to encrypt the data. Key groups and key records also are required for encryption.

Refer to the *NetBackup™ Security and Encryption Guide* for details.

Table 3-10 Required information for the encryption database

Field Name	Required information
KMS Server Name	This field displays the name of your NetBackup primary server. You can only configure KMS on your primary server. This field cannot be changed. If KMS is not configured, this field displays <code><kms_server_name></code> .
Host Master Key (HMK) Passphrase	Enter the key that protects the database. In KMS terminology, the key is called a <i>passphrase</i> .
Re-enter HMK Passphrase	Re-enter the host master key.
Host Master Key ID	The ID is a label that you assign to the master key. The ID lets you identify the particular host master key. You are limited to 255 characters in this field. To decipher the contents of a keystore file, you must identify the correct Key Protection Key and Host Master Key. These IDs are stored unencrypted in the keystore file header. You can select the correct ones even if you only have access to the keystore file. To perform a disaster recovery you must remember the correct IDs and the pass phrases that are associated with the files.
Key Protection Key (KPK) Passphrase	Enter the password that protects the individual records within the KMS database. In KMS terminology, the key is called a <i>passphrase</i> .
Re-enter KPK Passphrase	Re-enter the key protection password.

Table 3-10 Required information for the encryption database (*continued*)

Field Name	Required information
Key Protection Key ID	The ID is a label that you assign to the key. The ID lets you identify the particular key protection key. You are limited to 255 characters in this field. To decipher the contents of a keystore file, you must identify the correct Key Protection Key and Host Master Key. These IDs are stored unencrypted in the keystore file header. You can select the correct ones even if you only have access to the keystore file. To perform a disaster recovery you must remember the correct IDs and the pass phrases that are associated with the files.

After you configure the storage server and disk pool, it is recommended that you save a record of the key names.

See [“Saving a record of the KMS key names for NetBackup cloud storage encryption”](#) on page 114.

NetBackup cloud storage server properties

In **Host properties** you can change some of the properties that affect the NetBackup interaction with the cloud storage. The following table describes the prefixes that NetBackup uses to categorize the properties.

Not all properties apply to all storage vendors.

Table 3-11 Prefix definitions

Prefix	Definition	For more information
AMZ	Amazon	See “NetBackup cloud storage server connection properties” on page 106.
AMZGOV	Amazon GovCloud	See “NetBackup cloud storage server connection properties” on page 106.
AZR	Microsoft Azure	See “NetBackup cloud storage server connection properties” on page 106.
CLD	Cloudian Hyperstore	See “NetBackup cloud storage server connection properties” on page 106.
CRYPT	Encryption	See “NetBackup cloud storage server encryption properties” on page 112.
GOOG	Google Nearline	See “NetBackup cloud storage server connection properties” on page 106.

Table 3-11 Prefix definitions (*continued*)

Prefix	Definition	For more information
HT	Hitachi	See “NetBackup cloud storage server connection properties” on page 106.
HTTP	HTTP headers	See “NetBackup cloud storage server connection properties” on page 106. Note: This field applies to Amazon S3-compatible cloud providers.
METER	Metering	See “NetBackup cloud storage server connection properties” on page 106.
ORAC	Oracle Cloud	See “NetBackup cloud storage server connection properties” on page 106.
SWSTK-SWIFT	SwiftStack (Swift)	See “NetBackup cloud storage server connection properties” on page 106.
THR	Throttling	See “NetBackup cloud storage server bandwidth throttling properties” on page 102.
VER	Verizon	See “NetBackup cloud storage server connection properties” on page 106.

NetBackup cloud storage server bandwidth throttling properties

The following storage server properties apply to bandwidth throttling. The `THR` prefix specifies a throttling property. Use the correct cloud provider URL for the desired cloud vendor.

See [“Scalable Storage properties”](#) on page 74.

Table 3-12 Cloud storage server bandwidth throttling properties

Property	Description
THR:storage_server	Shows the maximum number of concurrent jobs that a specific cloud storage server can run. If you configure throttling for a media server that is a cloud storage server: ■ Change this value to 160 or more. ■ This value should be the same as the Maximum concurrent jobs media server property in the Scalable Storage host properties. See “Scalable Storage properties” on page 74. Default value: Not applicable Possible values: See the Description column
THR:AVAIL_BANDWIDTH	This read-only field displays the total available bandwidth value for the cloud feature. The value is displayed in bytes per second. You must specify a number greater than zero. If you enter zero, an error is generated. Default value: 104857600 Possible values: Any positive integer

Table 3-12 Cloud storage server bandwidth throttling properties (*continued*)

Property	Description
	The default maximum number of concurrent jobs that the media server can run for the cloud storage server. If <code>THR:storage_server</code> is set, NetBackup uses <code>THR:storage_server</code> instead of <code>THR:DEFAULT_MAX_CONNECTIONS</code>. This field is a read-only field. This value applies to the media server not to the cloud storage server. If you have more than one media server that can connect to the cloud storage server, each media server can have a different value. Therefore, to determine the total number of jobs that can run on the cloud storage server, add the values from each media server. If NetBackup is configured to allow more jobs than <code>THR:DEFAULT_MAX_CONNECTIONS</code>, NetBackup fails any jobs that start after the number of maximum jobs is reached. Jobs include both backup and restore jobs. You can configure job limits per backup policy and per storage unit. See the NetBackup Administrator's Guide, Volume I. Note: NetBackup must account for many factors when it starts jobs: the number of concurrent jobs, the number of <code>THR:DEFAULT_MAX_CONNECTIONS</code> per media server, the number of media servers, and the job load-balancing logic. Therefore, NetBackup may not fail jobs exactly at the maximum number of connections. NetBackup may fail a job when the connection number is slightly less than the maximum, exactly the maximum, or slightly more than the maximum. In practice, you should not need to set this value higher than 100. Default value: 10 Possible values: 1 to 2147483647
<code>THR:OFF_TIME_BANDWIDTH_PERCENT</code>	This read-only field displays the bandwidth percent that is used during off time. Default value: 100 Possible values: 0 to 100
<code>THR:OFF_TIME_END</code>	This read-only field displays the end of off time. Specify the time in 24-hour format. For example, 8:00 A.M. is 8 and 6:30 P.M. is 1830. Default value: 8 Possible values: 0 to 2359

Table 3-12 Cloud storage server bandwidth throttling properties (*continued*)

Property	Description
THR:OFF_TIME_START	This read-only field displays the start of off time. Specify the time in 24-hour format. For example, 8:00 A.M. is 8 and 6:30 P.M. is 1830. Default value: 18 Possible values: 0 to 2359
THR:READ_BANDWIDTH_PERCENT	This read-only field displays the read bandwidth percentage the cloud feature uses. Specify a value between 0 and 100. If you enter an incorrect value, an error is generated. Default value: 100 Possible values: 0 to 100
THR:SAMPLE_INTERVAL	This read-only field displays the rate at which backup streams sample their utilization and adjust their bandwidth use. The value is specified in seconds. When this value is set to zero, throttling is disabled. Default value: 0 Possible values: 1 to 2147483647
THR:WEEKEND_BANDWIDTH_PERCENT	This read-only field displays the bandwidth percent that is used during the weekend. Default value: 100 Possible values: 0 to 100
THR:WEEKEND_END	This read-only field displays the end of the weekend. The day value is specified with numbers, 1 for Monday, 2 for Tuesday, and so on. Default value: 7 Possible values: 1 to 7
THR:WEEKEND_START	This read-only field displays the start of the weekend. The day value is specified with numbers, 1 for Monday, 2 for Tuesday, and so on. Default value: 6 Possible values: 1 to 7
THR:WORK_TIME_BANDWIDTH_PERCENT	This read-only field displays the bandwidth percent that is used during the work time. Default value: 100 Possible values: 0 to 100

Table 3-12 Cloud storage server bandwidth throttling properties (*continued*)

Property	Description
THR:WORK_TIME_END	This read-only field displays the end of work time. Specify the time in 24-hour format. For example, 8:00 A.M. is 8 and 6:30 P.M. is 1830. Default value: 18 Possible values: 0 to 2359
THR:WORK_TIME_START	This read-only field displays the start of work time. Specify the time in 24-hour format. For example, 8:00 A.M. is 8 and 6:30 P.M. is 1830. Default value: 8 Possible values: 0 to 2359
THR:WRITE_BANDWIDTH_PERCENT	This read-only field displays the write bandwidth percentage the cloud feature uses. Specify a value between 0 and 100. If you enter an incorrect value, an error is generated. Default value: 100 Possible values: 0 to 100

NetBackup cloud storage server connection properties

All or most of the cloud storage servers use the storage server properties in [Table 3-13](#). The following are the prefixes for the currently supported cloud vendors:

- Amazon: AMZ
- Amazon GovCloud: AMZGOV
- Cloudian: CLD
- Google Nearline: GOOG
- Hitachi: HT
- Microsoft Azure: AZR
- Verizon: VER

Table 3-13 Storage server cloud connection properties

Property	Description
METER:DIRECTORY	This read-only field displays the directory in which to store data stream metering information. Default value: UNIX: <code>/usr/opensv/var/global/wmc/cloud</code> or <code>/usr/opensv/netbackup/db/cloud</code> (on media server versions 7.7.x to 8.1.2 only) Windows: <code>install_path\Veritas\NetBackup\var\global\wmc\cloud</code> or <code>install_path\Veritas\NetBackup\db\cloud\</code> (on media server versions 7.7.x to 8.1.2 only)
METER:INTERVAL	The interval at which NetBackup gathers connection information for reporting purposes. The value is set in seconds. The default setting is 300 seconds (5 minutes). If you set this value to zero, metering is disabled. To change this property, configure the Metering interval in the Scalable Storage host properties. See “Scalable Storage properties” on page 74. Default value: 300 Possible values: 1 to 10000
PREFIX:CURL_CONNECT_TIMEOUT	The amount of time that is allocated for the media server to connect to the cloud storage server. This value is specified in seconds. The default is 300 seconds or five minutes. This only limits the connection time, not the session time. If the media server cannot connect to the cloud storage server in the specified time, the job fails. This value cannot be disabled. If an invalid number is entered, the <code>CURL_CONNECT_TIMEOUT</code> returns to the default value of 300. Default value: 300 Possible values: 1 to 10000

Table 3-13 Storage server cloud connection properties (*continued*)

Property	Description
<code>PREFIX:CURL_TIMEOUT</code>	The maximum time in seconds to allow for the completion of a data operation. This value is specified in seconds. If the operation does not complete in the specified time, the operation fails. The default is 900 seconds (15 minutes). To disable this setting, set the value to 0 (zero). Default value: 900 Possible values: 1 to 10000
<code>PREFIX:LOG_CURL</code>	Determines if cURL activity is logged. The default is <code>NO</code> which means log activity is disabled. Default value: <code>NO</code> Possible values: <code>NO</code> (disabled) and <code>YES</code> (enabled)
<code>PREFIX:READ_BUFFER_SIZE</code>	The size of the buffer to use for read operations. <code>READ_BUFFER_SIZE</code> is specified in bytes. To enable the use of the buffer, set this value to a non-zero number. The <code>READ_BUFFER_SIZE</code> determines the size of the data packets that the storage server transmits during each restore job. An increase in the value may increase performance when a large amount of contiguous data is accessed. If insufficient bandwidth exists to transmit the specified amount of data within a few minutes, restore failures may occur due to time-outs. When you calculate the required bandwidth, consider the total load of simultaneous backup jobs and restore jobs on multiple media servers. See “About object size for cloud storage” on page 94.
<code>PREFIX:USE_SSL</code>	Determines if Secure Sockets Layer encryption is used for the control APIs. The default value is <code>YES</code>, meaning SSL is enabled. Default value: <code>YES</code> Possible values: <code>YES</code> or <code>NO</code>
<code>PREFIX:USE_SSL_RW</code>	Determines if Secure Sockets Layer encryption is used for read and write operations. The default value is <code>YES</code>, meaning SSL is enabled. Default value: <code>YES</code> Possible values: <code>YES</code> or <code>NO</code>

Table 3-13 Storage server cloud connection properties (*continued*)

Property	Description
<i>Provider Suffix:</i> USE_CRL	If SSL is enabled and the CRL option is enabled, each non-self-signed SSL certificate is verified against the CRL.
<i>PREFIX:</i> OBJECT_SIZE	The size of the data object that NetBackup sends to the cloud storage server with an HTTP PUT and GET requests. Object Size is specified in bytes. You cannot edit the Object Size once you set the value. See “About object size for cloud storage” on page 94.
<i>PREFIX:</i> WRITE_BUFFER_NUM	This parameter is not applicable for Amazon S3-compatible cloud providers. This read-only field displays the total number of write buffers the plug-in uses. The WRITE_BUFFER_SIZE value defines the size of the buffer. The value is set to 1 and cannot be changed. Default value: 1 Possible values: 1
<i>PREFIX:</i> WRITE_BUFFER_SIZE	The size of the buffer to use for write operations. WRITE_BUFFER_SIZE is specified in bytes. To disable the use of the buffer, set this value to 0 (zero). The WRITE_BUFFER_SIZE value determines the size of the data packs transmitted from the data mover to the storage server during a backup. An increase in the value may increase performance when a large amount of contiguous data is accessed. If insufficient bandwidth exists to transmit the specified amount of data within a few minutes, backup failures may occur due to time-outs. When you calculate the required bandwidth, consider the total load of simultaneous backup jobs and restore jobs on multiple media servers. See “About object size for cloud storage” on page 94.
HTTP:User-Agent	This property is applicable only for Amazon S3-compatible cloud providers. This property is set internally and the user cannot change this property.

Table 3-13 Storage server cloud connection properties (*continued*)

Property	Description
HTTP:x-amz-server-side-encryption	This property is applicable only for the following cloud providers: Amazon S3 and Amazon GovCloud Use this property to enable the server-side encryption of the data that you need to transfer to the cloud storage. AES-256 is a server-side encryption standard. Set this property to NONE to disable the server-side encryption for the cloud provider. Note: You should not enable this property if you already enabled the media server-side encryption option when you configured the cloud storage server.
AMZ:REGION_NAME	This property is applicable only for the Amazon GLACIER_VAULT storage class. Displays the region that is set during configuration of the storage server. This property is set during the configuration of the storage server and the user cannot change this property.
AMZ:UPLOAD_CLASS	This property is applicable only for the Amazon S3 Intelligent Tiering (LIFECYCLE) storage class. Use this property to specify the storage class to back up the data. Default value: STANDARD Possible values: STANDARD or STANDARD_IA
AMZ:RETRIEVAL RETENTION PERIOD	This property is applicable only for Amazon Glacier. Use this property to specify the retrieval retention period in days.
AMZ:TRANSITION_TO_STANDARD_IA_AFTER	This property is applicable only for the Amazon S3 Intelligent Tiering (LIFECYCLE) storage class. If you have set the UPLOAD_CLASS as STANDARD, the TRANSITION_TO_STANDARD_IA_AFTER must be set to either NONE or in the range 30 to 2147483617. If you have set the UPLOAD_CLASS as STANDARD_IA, the TRANSITION_TO_STANDARD_IA_AFTER must be set to NONE.

Table 3-13 Storage server cloud connection properties (*continued*)

Property	Description
AMZ:TRANSITION_TO_GLACIER_AFTER	This property is applicable only for the Amazon S3 Intelligent Tiering (LIFECYCLE) storage class. If you have set UPLOAD_CLASS as STANDARD, and if TRANSITION_TO_STANDARD_IA_AFTER is set in the range 30 to 2147483617, you must set TRANSITION_TO_GLACIER_AFTER as NONE or in the range 60 to 2147483647. This value includes a minimum stay of 30 days for the data in the STANDARD_IA storage class. If you have set UPLOAD_CLASS as STANDARD, and if TRANSITION_TO_STANDARD_IA_AFTER is set to NONE, you must set TRANSITION_TO_GLACIER_AFTER in the range 1 to 2147483647. If you have set UPLOAD_CLASS as STANDARD_IA and if TRANSITION_TO_STANDARD_IA_AFTER is set to NONE, you must set TRANSITION_TO_GLACIER_AFTER in the range 30 to 2147483647.
AMZ:STORAGE_CLASS	This property is applicable only for the Amazon S3 cloud providers. Displays the storage class that the cloud storage server uses. This property is set internally and the user cannot change this property.
AZR:STORAGE_TIER	This property is applicable only for Microsoft Azure Archive. Displays the storage tier that the cloud storage server uses.

Table 3-13 Storage server cloud connection properties (*continued*)

Property	Description
AMZ:OFFLINE_TRANSFER_MODE	This property is applicable only for the Amazon S3 cloud providers. Use this property to set the storage destination for Amazon Snowball. Default value: NONE Note: Set the property to NONE after you are done with using the Snowball mode. In this mode, the endpoint must point to Amazon public endpoint. Possible values: FILESYSTEM: Set this property if you want the data to be transferred to Amazon Snowball using the file interface. The storage server endpoint must point to the Amazon public endpoint. PROVIDER_API: Set this property if you want to transfer the data to Amazon Snowball using the S3 interface that Amazon provides. The storage server endpoint must point to Snowball endpoint.
AMZ:TRANSFER_DRIVE_PATH	This property is applicable only for the Amazon S3 cloud providers and if the AMZ:OFFLINE_TRANSFER_MODE property is set to FILESYSTEM Use this property to set the absolute mount point where the data must be backed up for Amazon Snowball. Default value: NONE

NetBackup cloud storage server encryption properties

The following encryption-specific storage server properties are used by all or most of the storage vendors. The `CRYPT` prefix specifies an encryption property. These values are for display purposes only and cannot be changed.

Table 3-14 Encryption cloud storage server properties

Property	Description
CRYPT:KMS_SERVER	This read-only field displays NetBackup server that hosts the KMS service. When you set the storage server properties, enter the name of the KMS server host. By default, this field contains the NetBackup primary server name. You cannot change this value. Default value: The NetBackup primary server name Possible values: N/A
CRYPT:KMS_VERSION	This read-only field displays the NetBackup Key Management Service version. You cannot change this value. Default value: 16 Possible values: N/A
CRYPT:LOG_VERBOSE	This read-only field displays if logs are enabled for encryption activities. The value is either YES for logging or NO for no logging. Default value: NO Possible values: YES and NO
CRYPT:VERSION	This read-only field displays the encryption version. You cannot change this value. Default value: 13107 Possible values: N/A

About cloud storage disk pools

A disk pool represents disk volumes on the underlying disk storage. A disk pool is the storage destination of a NetBackup storage unit. For cloud storage, you must specify only one volume for a disk pool.

Disk pool and disk volume names must be unique within your cloud storage provider's environment.

If a cloud storage disk pool is a storage destination in a storage lifecycle policy, NetBackup capacity management applies.

For instructions to create disk pools, see the *NetBackup™ Web UI Administrator's Guide*.

Configuring a disk pool for cloud storage

Use the web UI to create a disk pool for cloud storage. If you create encrypted storage and NetBackup KMS is configured, you must enter a pass phrase for each selected volume that uses encryption. The pass phrase creates the encryption key for that volume. If you create encrypted storage and external KMS is configured, you do not need to enter pass phrase for each selected volume.

See the *NetBackup™ Web UI Administrator's Guide* for the steps to create a disk pool.

Saving a record of the KMS key names for NetBackup cloud storage encryption

It is recommended that you save a record of the encryption key names and tags. The key tag is necessary if you need to recover or recreate the keys.

Saving a record of the NetBackup KMS server key names

Use the following procedure to save a record of the key names if NetBackup KMS server is configured when you enable the encryption setting during storage server configuration for cloud storage.

See [“About data encryption for cloud storage”](#) on page 91.

To save a record of the key names

- 1** To determine the key group names, use the following command on the primary server:

UNIX: `/usr/opensv/netbackup/bin/admincmd/nbkmsutil -listkgs`

Windows: `install_path\Program`

`Files\Veritas\NetBackup\bin\admincmd\nbkmsutil.exe -listkgs`

The following is example output:

```
Key Group Name       : CloudVendor.com:symc_backups_gold
Supported Cypher     : AES_256
Number of Keys       : 1
Has Active Key       : Yes
Creation Time        : Tues Oct 01 01:00:00 2013
Last Modification Time: Tues Oct 01 01:00:00 2013
Description          : CloudVendor.com:symc_backups_gold
```

Saving a record of the KMS key names for NetBackup cloud storage encryption

- 2 For each key group, write all of the keys that belong to the group to a file. Run the command on the primary server. The following is the command syntax:

UNIX: `/usr/openv/netbackup/bin/admincmd/nbkmsutil -listkeys -kgname
key_group_name > filename.txt`

Windows: `install_path\Program`

`Files\Veritas\NetBackup\bin\admincmd\nbkmsutil.exe -listkeys
-kgname key_group_name > filename.txt`

The following is example output:

```
nbkmsutil.exe -listkeys -kgname CloudVendor.com:symc_backups_gold
> encrypt_keys_CloudVendor.com_symc_backups_gold.txt
```

```
Key Group Name      : CloudVendor.com:symc_backups_gold
Supported Cypher    : AES_256
Number of Keys      : 1
Has Active Key      : Yes
Creation Time       : Tues Jan 01 01:00:00 2013
Last Modification Time: Tues Jan 01 01:00:00 2013
Description         : Key group to protect cloud volume
FIPS Approved Key   : Yes
```

```
Key Tag             : 532cf41cc8b3513a13c1c26b5128731e
                   : 5ca0b9b01e0689cc38ac2b7596bbae3c
Key Name            : Encrypt_Key_April
Current State       : Active
Creation Time       : Tues Jan 01 01:02:00 2013
Last Modification Time: Tues Jan 01 01:02:00 2013
Description         : -
```

```
Number of Keys: 1
```

- 3 Include in the file the pass phrase that you used to create the key record.
- 4 Store the file in a secure location.

Saving a record of an external KMS server key names

Refer to your KMS server documentation for key recovery steps.

Adding backup media servers to your cloud environment

You can add additional media servers to your cloud environment. Additional media servers can help improve backup performance. Such servers are known as *data movers*. The media servers that you add are assigned the credentials for the storage server. The credentials allow the data movers to communicate with the storage server.

A NetBackup media server must conform to the requirements for cloud storage.

See [“About the NetBackup media servers for cloud storage”](#) on page 97.

To add backup media servers to your cloud environment

- 1 On the left, click **Storage > Disk storage**.
- 2 Select the cloud storage server.
- 3 From the actions menu, select **Edit**.
- 4 Under **Media servers**, click **Add**.
- 5 Click **Save**.
- 6 Modify disk pools, storage units, and policies as desired.

Configuring a storage unit for cloud storage

Create one or more storage units that reference the disk pool.

To determine if storage units exist for the disk pool, click **Storage > Storage units** in the web UI.

A storage unit inherits the properties of the disk pool. If the storage unit inherits replication properties, the properties signal to a NetBackup storage lifecycle policy the intended purpose of the storage unit and the disk pool. Auto Image Replication requires storage lifecycle policies.

For details of configuring a storage unit, see *NetBackup™ Web UI Administrator's Guide*.

You can use storage unit properties to control your backup traffic.

See [“Configure a favorable client-to-server ratio”](#) on page 119.

See [“Control backup traffic to the media servers”](#) on page 120.

See [“Cloud storage unit properties”](#) on page 118.

Cloud storage unit properties

The following are the configuration options for a cloud disk pool storage unit.

Table 3-15 Cloud storage unit properties

Property	Description
Storage unit name	A unique name for the new storage unit. The name can describe the type of storage. The storage unit name is the name used to specify a storage unit for policies and schedules. The storage unit name cannot be changed after creation.
Storage unit type	Select Disk as the storage unit type.
Disk type	Select Cloud Storage (type) for the disk type. <i>type</i> represents the disk pool type, based on storage vendor, encryption, and so on.
Disk pool	Select the disk pool that contains the storage for this storage unit. All disk pools of the specified Disk type appear in the Disk pool list. If no disk pools are configured, no disk pools appear in the list.
Media server	The Media server setting specifies the NetBackup media servers that can backup clients and move the data to the cloud storage server. The media servers can also move the data for restore or duplication operations. Specify the media server or servers as follows: ■ To allow any server in the media server list to deduplicate data, select Use any available media server. ■ To use specific media servers to deduplicate the data, select Only use the following media servers. Then, select the media servers to allow. NetBackup selects the media server to use when the policy runs.

Table 3-15 Cloud storage unit properties (*continued*)

Property	Description
Maximum concurrent jobs	The Maximum concurrent jobs setting specifies the maximum number of jobs that NetBackup can send to a disk storage unit at one time. (Default: one job. The job count can range from 0 to 256.) This setting corresponds to the Maximum concurrent write drives setting for a Media Manager storage unit. NetBackup queues jobs until the storage unit is available. If three backup jobs are scheduled and Maximum concurrent jobs is set to two, NetBackup starts the first two jobs and queues the third job. If a job contains multiple copies, each copy applies toward the Maximum concurrent jobs count. Maximum concurrent jobs controls the traffic for backup and duplication jobs but not restore jobs. The count applies to all servers in the storage unit, not per server. If you select multiple media servers in the storage unit and 1 for Maximum concurrent jobs, only one job runs at a time. The number to enter depends on the available disk space and the server's ability to run multiple backup processes. Warning: A Maximum concurrent jobs setting of 0 disables the storage unit.
Maximum fragment size	For normal backups, NetBackup breaks each backup image into fragments so it does not exceed the maximum file size that the file system allows. You can enter a value from 20 MBs to 51200 MBs. For a FlashBackup policy, it is recommended that you use the default, maximum fragment size to ensure optimal duplication performance.

Configure a favorable client-to-server ratio

You can use storage unit settings to configure a favorable client-to-server ratio. You can use one disk pool and configure multiple storage units to separate your backup traffic. Because all storage units use the same disk pool, you do not have to partition the storage.

For example, assume that you have 100 important clients, 500 regular clients, and four media servers. You can use two media servers to back up your most important clients and two media servers to back up your regular clients.

The following example describes how to configure a favorable client-to-server ratio:

- Configure the media servers for NetBackup deduplication and configure the storage.

- Configure a disk pool.
- Configure a storage unit for your most important clients (such as STU-GOLD). Select the disk pool. Select **Only use the following media servers**. Select two media servers to use for your important backups.
- Create a backup policy for the 100 important clients and select the STU-GOLD storage unit. The media servers that are specified in the storage unit move the client data to the deduplication storage server.
- Configure another storage unit (such as STU-SILVER). Select the same disk pool. Select **Only use the following media servers**. Select the other two media servers.
- Configure a backup policy for the 500 regular clients and select the STU-SILVER storage unit. The media servers that are specified in the storage unit move the client data to the deduplication storage server.

Backup traffic is routed to the wanted data movers by the storage unit settings.

Note: NetBackup uses storage units for media server selection for write activity (backups and duplications) only. For restores, NetBackup chooses among all media servers that can access the disk pool.

Control backup traffic to the media servers

On disk pool storage units, you can use the **Maximum concurrent jobs** settings to control the backup traffic to the media servers. Effectively, this setting directs higher loads to specific media servers when you use multiple storage units for the same disk pool. A higher number of concurrent jobs means that the disk can be busier than if the number is lower.

For example, two storage units use the same set of media servers. One of the storage units (STU-GOLD) has a higher **Maximum concurrent jobs** setting than the other (STU-SILVER). More client backups occur for the storage unit with the higher **Maximum concurrent jobs** setting.

About NetBackup Accelerator and NetBackup Optimized Synthetic backups

NetBackup Cloud Storage supports NetBackup Accelerator and NetBackup Optimized Synthetics. Encryption, metering, and throttling are functional and supported when you enable NetBackup Accelerator or NetBackup Optimized Synthetic backups. You enable both NetBackup Accelerator and NetBackup

Optimized Synthetic backups in the same way as non-Cloud backups. More information about NetBackup Accelerator and NetBackup Optimized Synthetic backups is available.

- See the [NetBackup Deduplication Guide](#).
- See the [NetBackup Administrator's Guide, Volume I](#)

Cloud storage disk pool properties

The properties of a disk pool may vary depending on the purpose the disk pool.

Note: The following properties do not apply to cloud storage disk pools: **Total available space**, **Total raw size**, **Usable Size**, **Low water mark**, and **High water mark**.

All these values are derived from the storage capacity, which cannot be fetched from the cloud provider.

The following table describes the possible properties:

Table 3-16 Cloud storage disk pool properties

Property	Description
Name	The disk pool name.
Storage servers	The storage server name.
Disk volumes	The disk volume that comprises the disk pool.
Total raw size	The total raw, unformatted size of the storage in the disk pool. The storage host may or may not expose the raw size of the storage. Note: Total raw size does not apply to cloud storage disk pools.
Total available space	The total amount of space available in the disk pool. Note: Total available space does not apply to cloud storage disk pools.
Comments	A comment that is associated with the disk pool.

Table 3-16 Cloud storage disk pool properties (*continued*)

Property	Description
High water mark	The High water mark, is a threshold at which the volume or the disk pool is considered full. Note: High water mark does not apply to cloud storage disk pools.
Low water mark	The Low water mark is a threshold at which NetBackup stops image cleanup. Low water mark does not apply to cloud storage disk pools.
Limit I/O streams	Select to limit the number of read and write streams (that is, jobs) for each volume in the disk pool. A job may read backup images or write backup images. By default, there is no limit. When the limit is reached, NetBackup chooses another volume for write operations, if available. If not available, NetBackup queues jobs until a volume is available. Too many streams may degrade performance because of disk thrashing. Disk thrashing is excessive swapping of data between RAM and a hard disk drive. Fewer streams can improve throughput, which may increase the number of jobs that complete in a specific time period. A starting point is to divide the Maximum concurrent jobs of all of the storage units by the number of volumes in the disk pool.
per volume	Select or enter the number of read and write streams to allow per volume. Many factors affect the optimal number of streams. Factors include but are not limited to disk speed, CPU speed, and the amount of memory. For the disk pools that are configured for Snapshot and that have a Replication source property: ■ Always use increments of 2 when you change this setting. A single replication job uses two I/O streams. ■ If more replication jobs exist than streams are available, NetBackup queues the jobs until streams are available. ■ Batching can cause many replications to occur within a single NetBackup job. Another setting affects snapshot replication job batching.

Certificate validation against Certificate Revocation List (CRL)

For all the cloud providers, NetBackup provides a capability to verify the SSL certificates against the CRL (Certificate Revocation List). If SSL is enabled and the CRL option is enabled, each non-self-signed SSL certificate is verified against the CRL. If the certificate is revoked, NetBackup does not connect to the cloud provider.

You can enable validation against CRL using one of the following ways:

- `csconfig` CLI: `crl` parameter is added to the SSL parameters. The option is available when you add or update the storage server. You can change the CRL value only through the `csconfig` CLI, before you create an alias.
- Storage server properties dialog: Update the `USE_CRL` property from the storage server properties dialog. From the GUI, you can only disable the CRL option, after configuration.
See [“NetBackup cloud storage server connection properties”](#) on page 106.
- You can also use to the `nbdevconfig` CLI with `getconfig` and `setconfig` options to enable or disable verification against CRL.

Note: Post upgrade, for the cloud storage servers with SSL enabled, the CRL validation is enabled by default.

Requirements for enabling certificate validation against Certificate Revocation List (CRL)

- CRL distribution endpoints are HTTP thus, turn off any firewall rule that block HTTP (port 80) connection to external network. For example, `http://crl3.provider.com/server-g2.crl`
- CRL download URL is dynamically fetched from the certificate thus, disable any firewall rule that blocks unknown URLs.
- Typically, CRL URLs (distribution endpoints) support IPV4. For IPV6 environments disable the CRL option.
- Private Clouds typically have a self-signed certificate. Thus, for private clouds, CRL check is not required. The check is skipped even if CRL option is enabled.
- CRL distribution point must be present in the x.509 certificate. The type of distribution point must HTTP.

Managing Certification Authorities (CA) for NetBackup Cloud

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

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

- Windows:
 - On media server versions 10.0 and later, the path is:
`<installation-path>\NetBackup\var\global\cloud`
 - On media server versions 8.2 to 9.1, the path is:
`<installation-path>\NetBackup\var\global\wmc\cloud\cacert.pem`
 - On media server versions 7.7.x to 8.1.2, the path is:
`install_path\Veritas\NetBackup\db\cloud\cacert.pem.`
- UNIX:
 - On media server versions 10.0 and later, the path is:
`/usr/opensv/var/global/cloud/`
 - On media server versions 8.2 to 9.1, the path is:
`/usr/opensv/var/global/wmc/cloud/cacert.pem`
 - On media server versions 7.7.x to 8.1.2, the path is:
`/usr/opensv/netbackup/db/cloud/cacert.pem.`

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

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

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

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

To add a CA

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

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

Add the following information block:

```
Certificate Authority Name
=====

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

<Certificate content>

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

- 3 Save the file.

To remove a CA

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

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

```
Certificate Authority Name
=====

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

<Certificate content>

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

- 3 Save the file.

List of CAs approved by NetBackup

- AddTrust External Root
- Baltimore CyberTrust Root
- Cybertrust Global 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
- GeoTrust Global CA
- GeoTrust Primary Certification Authority
- GeoTrust Primary Certification Authority - G2
- GeoTrust Primary Certification Authority - G3
- GeoTrust Universal CA
- GeoTrust Universal CA 2
- RSA Security 2048 v3
- Starfield Services Root Certificate Authority - G2
- Thawte Primary Root CA
- Thawte Primary Root CA - G2
- Thawte Primary Root CA - G3
- VeriSign Class 1 Public Primary Certification Authority - G3
- VeriSign Class 2 Public Primary Certification Authority - G3
- Verisign Class 3 Public Primary Certification Authority - G3
- VeriSign Class 3 Public Primary Certification Authority - G4
- VeriSign Class 3 Public Primary Certification Authority - G5
- VeriSign Universal Root Certification Authority

Monitoring and Reporting

This chapter includes the following topics:

- [Viewing cloud storage job details](#)
- [Viewing the compression ratio](#)
- [Displaying KMS key information for cloud storage encryption](#)

Viewing cloud storage job details

Use the NetBackup Activity Monitor to view job details.

To view cloud storage job details

- 1 On the left, click **Activity Monitor**.
- 2 Click the **Jobs** tab.
- 3 To view the details for a specific job, click on the job that is displayed in the **Jobs** table.

Viewing the compression ratio

The bptm logs provide information of the compression ratio of your data after the backup is taken in the cloud storage. The compression ratio is calculated by dividing the original size with the compressed size. For example, if the original data is of 15302918144 bytes and is compressed to 7651459072, then the compression ratio is 2.00.

To view the compression ratio

- 1 Note down the bptm PID of the backup job.
See [“Viewing cloud storage job details”](#) on page 127.
- 2 Open the `bptm.log` file. The log file resides in the following directories:

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

Windows `install_path\NetBackup\logs\`

- 3 Search for the bptm PID instance.

The following lines provide the compression ratio information according to the image format:

```
date:time <PID> <4> 35:bptm:<PID>:  
media_server_IP: compress: image image_name_C1_F1 compressed from data in
```

```
date:time <PID> <4> 35:bptm:<PID>:  
media_server_IP: compress: image image_name_C1_HDR compressed from data i
```

Displaying KMS key information for cloud storage encryption

You can use the `nbkmsutil` command to list the following information about the key groups and the key records:

Key groups See [To display KMS key group information](#).

Keys See [To display KMS key information](#).

Note: It is recommended that you keep a record key information. The key tag that is listed in the output is necessary if you need to recover keys.

To display KMS key group information

- ◆ To list all of the key groups, use the `nbkmsutil` with the `-listkgs` option. The following is the command format:

UNIX: `/usr/opensv/netbackup/bin/admincmd/nbkmsutil -listkgs`

Windows: `install_path\Veritas\NetBackup\bin\admincmd\nbkmsutil -listkgs`

The following is example output on UNIX hosted storage. On Windows, the volume name is not used.

```
nbkmsutil -listkgs
```

```
Key Group Name      : CloudStorageVendor.com:symc_volume_for_backups
Supported Cypher    : AES_256
Number of Keys      : 1
Has Active Key      : Yes
Creation Time       : Tues Jan 01 01:00:00 2013
Last Modification Time: Tues Jan 01 01:00:00 2013
Description         : -
```

To display KMS key information

- ◆ To list all of the keys that belong to a key group name, use the `nbkmsutil` with the `-listkgs` and `-kgname` options. The following is the command format:

UNIX: `/usr/openv/netbackup/bin/admincmd/nbkmsutil -listkeys -kgname AdvDiskServer1.example.com:AdvDisk_Volume`

Windows: `install_path\Veritas\NetBackup\bin\admincmd\nbkmsutil -listkeys -kgname AdvDiskServer1.example.com:`

The following is example output on UNIX hosted storage. On Windows, the volume name is not used.

```
nbkmsutil -listkeys -kgname CloudStorageVendor.com:symc_volume_for_backup
```

```
Key Group Name      : CloudStorageVendor.com:symc_volume_for_backups
Supported Cypher    : AES_256
Number of Keys      : 1
Has Active Key      : Yes
Creation Time       : Tues Jan 01 01:00:00 2013
Last Modification Time: Tues Jan 01 01:00:00 2013
Description         : -
```

```
Key Tag            : 532cf41cc8b3513a13c1c26b5128731e5ca0b9b01e0689cc38ac2b7596bbae3c
Key Name           : Encrypt_Key_April
Current State      : Active
Creation Time      : Tues Jan 01 01:02:00 2013
Last Modification Time: Tues Jan 01 01:02:00 2013
Description        : -
```

You can also use the `nbkmscmd` command to list the keys from NetBackup KMS and external KMS server. You need to ensure that a Symmetric encryption key already exists in the external KMS server with a custom attribute with value of key group in the 'storage_server_name:volume_name' format.

To display the key information for NetBackup KMS and external KMS

- 1 Run the following command to retrieve the KMS server configuration names.

```
nbkmscmd -listkmsconfig
```

- 2 Run the following command to retrieve key information for a key group from the KMS server.

```
nbkmscmd -listkeys -name KMS_server_name -keyGroupName
key_group_name -jsonRaw
```

Operational notes

This chapter includes the following topics:

- [NetBackup bpstsinfo command operational notes](#)
- [Cloud configuration may fail if NetBackup Access Control is enabled](#)
- [Deleting cloud storage server artifacts](#)
- [Using csconfig reinitialize to load updated cloud configuration settings](#)
- [Enabling or disabling communication between primary server and legacy cloud storage media servers](#)

NetBackup bpstsinfo command operational notes

The following table describes operational notes for the `bpstsinfo` command with NetBackup cloud storage.

Table 5-1 `bpstsinfo` command operational notes

Note	Description
Use either the <code>-stype</code> option or the <code>-storageserverprefix</code>	Use either the <code>-stype</code> option or the <code>-storageserverprefix</code> option to constrain the <code>bpstsinfo</code> command to list storage server information. If you do not, the command searches all providers, which may be time consuming and may result in a timeout.
Specify the correct <code>-stype</code>	The plug-in that requests the information affects the information that is returned. Therefore, use the correct <code>-stype</code> with the <code>bpstsinfo</code> command. To determine the <code>-stype</code>, use the following command: <pre>nbdevquery -liststs -storage_server fq_host_name</pre> If the storage is encrypted, the <code>-stype</code> includes an <code>_crypt</code> suffix.

Table 5-1 `bpstsinfo` command operational notes (*continued*)

Note	Description
Encrypted and non-encrypted storage units are displayed in <code>bpstsinfo</code> command output	When you use the <code>bpstsinfo</code> command to display the encrypted logical storage unit (LSU) information, the output shows both encrypted and non-encrypted LSUs if both types exist. That output is the expected result. The <code>bpstsinfo</code> command operates on the level of the storage plug-in, which is not aware of any higher-level detail, such as encryption. The following is an example of a command that specifies encrypted storage: <pre>bpstsinfo -lsuinfo -storage_server amazon.com -stype amazon_crypt</pre>

Cloud configuration may fail if NetBackup Access Control is enabled

If you attempt to configure a cloud storage server in an environment that uses NetBackup Access Control, you may receive an error message similar to the following:

```
Error creating Key Group and Keys cannot connect on socket
```

NetBackup generates this error message because the user does not have sufficient rights within NetBackup Access Control. The user account that configures the cloud storage server must be a member of the NBU_KMS Admin Group.

See the [NetBackup Security and Encryption Guide](#) for more information about NetBackup Access Control and account setup:

Deleting cloud storage server artifacts

If you incorrectly remove a storage server, configuration files are left orphaned on the computer. Attempts to create a new storage server fail with an error message that indicates a logon failure. Use the following procedure to correctly delete a storage server:

Deleting a storage server

- 1 Expire all images on the storage server.
- 2 Delete the storage unit.
- 3 Delete the disk pool.

- 4 Delete the storage server.
- 5 Delete .pref files from db/cloud directory.

Using csconfig reinitialize to load updated cloud configuration settings

You might update your NetBackup cloud storage configuration settings, typically when you have upgraded the NetBackup primary server or have downloaded a newer version of the NetBackup Cloud configuration package (CloudProvider.xml configuration file). When you install the updated package or make updates to your existing cloud storage configuration settings, then depending on the NetBackup release version, you are required to restart the NetBackup CloudStore Service Container (nbcssc) or the NetBackup Web Management Console (nbwmc) service for the configuration changes to take effect.

Sometimes, the nbcssc or nbwmc service might hang and a service restart might fail. This happens either due to an invalid CloudProvider.xml file or due to a version mismatch between the xml file and the configured CloudStore version. A service restart failure can eventually lead to a failure in the NetBackup backup jobs.

Starting with NetBackup 8.2 release, you can use the csconfig utility to reload the updated cloud configuration settings without the need to restart any service.

After making the configuration updates, run the following command on the NetBackup primary or media server:

On UNIX, run the following command from the
/usr/opensv/netbackup/bin/admincmd/ directory:

```
# sudo ./csconfig reinitialize
```

On Windows, run the following command from the
<install_path>\NetBackup\bin\admincmd\ directory:

```
csconfig reinitialize
```

When you run the csconfig reinitialize command option, the nbwmc service reloads the configuration settings from the Cloudstore.conf, CloudProvider.xml, and CloudInstance.xml files. There is no need to restart the nbwmc service.

Enabling or disabling communication between primary server and legacy cloud storage media servers

This is applicable for media server version 7.7.x to 8.1.2 only.

The NetBackup CloudStore Service Container (`nbcssc`) service that runs on older cloud storage media servers uses port 5637 to communicate with the primary server. Starting with release 8.2, `nbcssc` service is no longer deployed. The NetBackup Web Management Console (`nbwmc`) and the NetBackup Service Layer (`nbsl`) services now handle that functionality.

Even when you upgrade your primary server to 8.2 or later, the legacy cloud storage media servers continue to use the legacy cloud service for communicating with the primary server. The NetBackup 8.2 primary server, however, does support legacy cloud storage media servers. To allow communication between an 8.2 primary server and the older media servers, you have to open port 5637 on the primary server.

To enable nbwmc service communication on port 5637

- 1 Run the following command on the primary server:

UNIX:

```
# usr/opensv/wmc/bin/install/configurePorts -addLegacyCloudService
```

Windows:

```
<install_path>\NetBackup\wmc\bin\install\configurePorts
-addLegacyCloudService
```

- 2 Restart the `nbwmc` service for the changes to take effect.

- 3 Run the following command to provision a hostname-based certificate for the media server:

UNIX:

```
# usr/opensv/netbackup/bin/admincmd/bpnbaz -ProvisionCert
<media_server>
```

Windows:

```
<install_path>\NetBackup\bin\admincmd\bpnbaz -ProvisionCert
<media_server>
```

In case of an appliance, run the following commands instead:

UNIX:

```
# usr/opensv/netbackup/bin/bpnbat -AddMachine <appliance_hostname>
```

Windows:

```
<install_path>\NetBackup\bin\bpnbat -AddMachine
<appliance_hostname>
```

- 4 Restart the cloud storage services on the media server.

Even though older versions of media servers are supported, it is recommended that you upgrade such media servers to version 8.2 or later. After upgrading all the legacy media servers, you can disable `nbwmc` service usage on port 5637.

To disable nbwmc service communication on port 5637

- 1 Run the following command on the primary server:

UNIX:

```
# usr/opensv/wmc/bin/install/configurePorts
-removeLegacyCloudService
```

Windows:

```
<install_path>\NetBackup\wmc\bin\install\configurePorts
-removeLegacyCloudService
```

- 2 Restart the `nbwmc` service for the changes to take effect.

Troubleshooting

This chapter includes the following topics:

- [About unified logging](#)
- [About legacy logging](#)
- [NetBackup cloud storage log files](#)
- [Enable libcurl logging](#)
- [Troubleshooting cloud storage configuration issues](#)
- [Troubleshooting cloud storage operational issues](#)
- [Troubleshooting Amazon Snowball and Amazon Snowball Edge issues](#)

About unified logging

Unified logging creates log file names and messages in a format that is standardized across Cohesity products. Only the `vxlogview` command can assemble and display the log information correctly. Server processes and client processes use unified logging.

Log files for originator IDs are written to a subdirectory with the name specified in the log configuration file. All unified logs are written to subdirectories in the following directory:

Windows `install_path\NetBackup\logs`

UNIX `/usr/opensv/logs`

Note: Only the following types of users can access the logs: root and service users in Linux systems, and users present in the administrators group of Windows systems.

You can access logging controls in **Logging** host properties. You can also manage unified logging with the following commands:

<code>vxlogcfg</code>	Modifies the unified logging configuration settings.
<code>vxlogmgr</code>	Manages the log files that the products that support unified logging generate.
<code>vxlogview</code>	Displays the logs that unified logging generates.

See [“Examples of using vxlogview to view unified logs”](#) on page 139.

About using the `vxlogview` command to view unified logs

Only the `vxlogview` command can assemble and display the unified logging information correctly. The unified logging files are in binary format and some of the information is contained in an associated resource file. These logs are stored in the following directory. You can display `vxlogview` results faster by restricting the search to the files of a specific process.

UNIX	<code>/usr/opensv/logs</code>
Windows	<code>install_path\NetBackup\logs</code>

Table 6-1 Fields in `vxlogview` query strings

Field name	Type	Description	Example
PRODID	Integer or string	Provide the product ID or the abbreviated name of product.	PRODID = 51216 PRODID = 'NBU'
ORGID	Integer or string	Provide the originator ID or the abbreviated name of the component.	ORGID = 116 ORGID = 'nbpem'
PID	Long Integer	Provide the process ID	PID = 1234567
TID	Long Integer	Provide the thread ID	TID = 2874950
STDATE	Long Integer or string	Provide the start date in seconds or in the locale-specific short date and time format. For example, a locale can have the format 'mm/dd/yy hh:mm:ss AM/PM'	STDATE = 98736352 STDATE = '4/26/11 11:01:00 AM'

Table 6-1 Fields in vxlogview query strings (*continued*)

Field name	Type	Description	Example
ENDATE	Long Integer or string	Provide the end date in seconds or in the locale-specific short date and time format. For example, a locale can have the format 'mm/dd/yy hh:mm:ss AM/PM'	ENDATE = 99736352 ENDATE = '04/27/11 10:01:00 AM'
PREVTIME	String	Provide the hours in 'hh:mm:ss' format. This field should be used only with operators =, <, >, >=, and <=	PREVTIME = '2:34:00'
SEV	Integer	Provide one of the following possible severity types: 0 = INFO 1 = WARNING 2 = ERR 3 = CRIT 4 = EMERG	SEV = 0 SEV = INFO
MSGTYPE	Integer	Provide one of the following possible message types: 0 = DEBUG (debug messages) 1 = DIAG (diagnostic messages) 2 = APP (application messages) 3 = CTX (context messages) 4 = AUDIT (audit messages)	MSGTYPE = 1 MSGTYPE = DIAG
CTX	Integer or string	Provide the context token as string identifier or 'ALL' to get all the context instances to be displayed. This field should be used only with the operators = and !=.	CTX = 78 CTX = 'ALL'

Table 6-2 Examples of query strings with dates

Example	Description
<code>(PRODID == 51216) && ((PID == 178964) ((STDATE == '2/5/15 09:00:00 AM') && (ENDATE == '2/5/15 12:00:00 PM')))</code>	Retrieves the log file message for the NetBackup product ID 51216 between 9AM and 12PM on 2015-05-02.
<code>((prodid = 'NBU') && ((stdate >= '11/18/14 00:00:00 AM') && (enddate <= '12/13/14 12:00:00 PM')))) ((prodid = 'BENT') && ((stdate >= '12/12/14 00:00:00 AM') && (enddate <= '12/25/14 12:00:00 PM'))))</code>	Retrieves the log messages for the NetBackup product NBU between 2014-18-11 and 2014-13-12 and the log messages for the NetBackup product BENT between 2014-12-12 and 2014-25-12.
<code>(STDATE <= '04/05/15 0:0:0 AM')</code>	Retrieves the log messages that were logged on or before 2015-05-04 for all of the installed Cohesity products.

Examples of using vxlogview to view unified logs

The following examples demonstrate how to use the `vxlogview` command to view unified logs.

Note: Only the following types of users can access the logs: root and service users in Linux systems, and users present in the administrators group of Windows systems.

Table 6-3 Example uses of the vxlogview command

Item	Example
Display all the attributes of the log messages	<code>vxlogview -p 51216 -d all</code>
Display specific attributes of the log messages	Display the log messages for NetBackup (51216) that show only the date, time, message type, and message text: <code>vxlogview --prodid 51216 --display D,T,m,x</code>
Display the latest log messages	Display the log messages for originator 116 (nbpem) that were issued during the last 20 minutes. Note that you can specify <code>-o nbpem</code> instead of <code>-o 116</code> : <code># vxlogview -o 116 -t 00:20:00</code>

Table 6-3 Example uses of the vxlogview command (*continued*)

Item	Example
Display the log messages from a specific time period	Display the log messages for nbpem that were issued during the specified time period: <pre># vxlogview -o nbpem -b "05/03/15 06:51:48 AM" -e "05/03/15 06:52:48 AM"</pre>
Display results faster	You can use the <code>-i</code> option to specify an originator for a process: <pre># vxlogview -i nbpem</pre> The <code>vxlogview -i</code> option searches only the log files that the specified process (nbpem) creates. By limiting the log files that it has to search, <code>vxlogview</code> returns a result faster. By comparison, the <code>vxlogview -o</code> option searches all unified log files for the messages that the specified process has logged. Note: If you use the <code>-i</code> option with a process that is not a service, <code>vxlogview</code> returns the message "No log files found." A process that is not a service has no originator ID in the file name. In this case, use the <code>-o</code> option instead of the <code>-i</code> option. The <code>-i</code> option displays entries for all OIDs that are part of that process including libraries (137, 156, 309, etc.).
Search for a job ID	You can search the logs for a particular job ID: <pre># vxlogview -i nbpem grep "jobid=job_ID"</pre> The <code>jobid=</code> search key should contain no spaces and must be lowercase. When searching for a job ID, you can use any <code>vxlogview</code> command option. This example uses the <code>-i</code> option with the name of the process (nbpem). The command returns only the log entries that contain the job ID. It misses related entries for the job that do not explicitly contain the <code>jobid=job_ID</code>.

About legacy logging

In NetBackup legacy debug logging, a process creates log files of debug activity in its own logging directory. By default, NetBackup creates only a subset of logging directories, in the following locations:

Windows	<code>install_path\NetBackup\logs</code> <code>install_path\Volmgr\debug</code>
UNIX	<code>/usr/opensv/netbackup/logs</code> <code>/usr/opensv/volmgr/debug</code>

To use legacy logging, a log file directory must exist for a process. If the directory is not created by default, you can use the `mklogdir` utility to create the directories. Or, you can manually create the directories. When logging is enabled for a process, a log file is created when the process begins. Each log file grows to a certain size before the NetBackup process closes it and creates a new log file.

Note: To apply the appropriate permissions on the legacy log directories, always use the `mklogdir` utility present in Windows and Linux to create the legacy log directories for each platform.

You can use the following utility to create all of the log directories:

- **Windows:** `install_path\NetBackup\Logs\mklogdir.bat`
- **UNIX:** `/usr/opensv/netbackup/logs/mklogdir`

Follow these recommendations when you create and use legacy log folders:

- Do not use symbolic links or hard links inside legacy log folders.
- Sometimes if a process runs for a non-root or non-admin user, no logging that occurs in the legacy log folders. In that case use the `mklogdir` command to create a folder for the required user.
- To run a command line for a non-root or non-admin user (troubleshooting when the NetBackup services are not running), create user folders for the specific command line. Create the folders either with the `mklogdir` command or manually with the non-root or non-admin user privileges.

Creating NetBackup log file directories for cloud storage

Before you configure your NetBackup feature, create the directories into which the NetBackup commands write log files. Create the directories on the primary server and on each media server that you use for your feature. The log files reside in the following directories:

- **UNIX:** `/usr/opensv/netbackup/logs/`
- **Windows:** `install_path\NetBackup\logs\`

More information about NetBackup logging is available in the [NetBackup Logging Reference Guide](#).

To create log directories for NetBackup commands

- ◆ Depending on the operating system, run one of the following scripts:

UNIX: `/usr/opensv/netbackup/logs/mklogdir`

Windows: `install_path\NetBackup\logs\mklogdir.bat`

To create the `tpconfig` command log directory

- ◆ Depending on the operating system, create the `debug` directory and the `tpcommand` directory (by default, the `debug` directory and the `tpcommand` directory do not exist). The pathnames of the directories are as follows:

UNIX: `/usr/opensv/volmgr/debug/tpcommand`

Windows: `install_path\Veritas\Volmgr\debug\tpcommand`

NetBackup cloud storage log files

NetBackup cloud storage exists within the Cohesity OpenStorage framework. Therefore, the log files for cloud activity are the same as for OpenStorage with several additions.

Some NetBackup commands or processes write messages to their own log files. For those commands and processes, the log directories must exist so that the utility can write log messages.

Other processes use Veritas Unified Logging (VxUL). Each process has a corresponding VxUL originator ID. VxUL uses a standardized name and file format for log files. To view VxUL log files, you must use the NetBackup `vxlogview` command.

More information about how to view and manage log files is available. See the [NetBackup Logging Reference Guide](#).

The following are the component identifiers for log messages:

- An `sts_` prefix relates to the interaction with the plug-in that writes to and reads from the storage.
- A cloud storage server prefix relates to interaction with that cloud vendor's storage network.
- An `encrypt` prefix relates to interaction with the encryption plug-in.
- A `KMSCLIB` prefix relates to interaction with the NetBackup Key Management Service.

Most interaction occurs on the NetBackup media servers. Therefore, the log files on the media servers that you use for disk operations are of most interest.

Warning: The higher the log level, the greater the affect on NetBackup performance. Use a log level of 5 (the highest) only when directed to do so by a Cohesity representative. A log level of 5 is for troubleshooting only.

Specify the NetBackup log levels in the **Logging** host properties on the NetBackup primary server. The log levels for some processes specific to certain options are set in configuration files as described in [Table 6-4](#).

[Table 6-4](#) describes the logs.

Table 6-4 NetBackup logs for cloud storage

Activity	OID	Processes
Backups and restores	N/A	Messages appear in the log files for the following processes: ■ The <code>bpbrm</code> backup and restore manager. ■ The <code>bpdbm</code> database manager. ■ The <code>bpdm</code> disk manager. ■ The <code>bptm</code> tape manager for I/O operations. The log files reside in the following directories: ■ UNIX: <code>/usr/opensv/netbackup/logs/</code> ■ Windows: <code>install_path\NetBackup\logs\</code>
Backups and restores	117	The <code>nbjm</code> Job Manager.
Image cleanup, verification, import, and duplication	N/A	The <code>bpdbm</code> database manager log files. The log files reside in the following directories: ■ UNIX: <code>/usr/opensv/netbackup/logs/bpdbm</code> ■ Windows: <code>install_path\NetBackup\logs\bpdbm</code>
Cloud connection operations	N/A	The <code>bpstsinfo</code> utility writes information about connections to the cloud storage server in its log files.
Cloud account configuration	222	The Remote Manager and Monitor Service is the process that creates the cloud storage accounts. RMMS runs on media servers.
Cloud Storage Service Container	N/A	This is applicable to media server versions 7.7.x to 8.1.2 only. The NetBackup Cloud Storage Service Container (<code>nbcssc</code>) writes log files to the following directories: ■ For Windows: <code>install_path\Veritas\NetBackup\logs\nbcssc</code> ■ For UNIX: <code>/usr/opensv/netbackup/logs/nbcssc</code>

Table 6-4 NetBackup logs for cloud storage (*continued*)

Activity	OID	Processes
NetBackup Web Management Console	495	The NetBackup Web Management Console (<code>nbwmc</code>) service writes logs to the following directories: - For Windows: <code>install_path\Veritas\netbackup\logs\nbweb service</code> - For UNIX: <code>/usr/openv/logs/nbweb service</code>
NetBackup Service Layer	N/A	The NetBackup Service Layer (<code>nbsl</code>) service writes logs to the following directories: - For Windows: <code>install_path\Veritas\netbackup\logs\nbsl</code> - For UNIX: <code>/usr/openv/logs/nbsl</code>
csconfig utility	N/A	The NetBackup <code>csconfig</code> command-line utility writes logs to the following directories: - For Windows: <code>install_path\Veritas\netbackup\logs\ncssc</code> - For UNIX: <code>/usr/openv/netbackup/logs/ncssc</code>
Credentials configuration	N/A	The <code>tpconfig</code> utility. The <code>tpconfig</code> command writes log files to the <code>tpcommand</code> directory.
Device configuration	111	The <code>nbemm</code> process.
Device configuration	178	The Disk Service Manager process that runs in the Enterprise Media Manager (EMM) process.
Device configuration	202	The Storage Server Interface process that runs in the Remote Manager and Monitor Service. RMMS runs on media servers.
Device configuration	230	The Remote Disk Service Manager interface (RDSM) that runs in the Remote Manager and Monitor Service. RMMS runs on media servers.

See [“Troubleshooting cloud storage operational issues”](#) on page 149.

Enable libcurl logging

Set the storage server property `CLOUD_PREFIX:LOG_CURL` to `YES` to enable cURL logging. The `CLOUD_PREFIX` value is the prefix value of each storage provider. The possible values are:

AMZ	Amazon
AMZGOV	Amazon GovCloud
AZR	Microsoft Azure
CLD	Cloudian HyperStore
GOOG	Google Nearline
HT	Hitachi
ORAC	Oracle Cloud
SWSTK-SWIFT	SwiftStack (Swift)
VER	Verizon

For example, to enable `LOG_CURL` for Amazon, set `AMZ:LOG_CURL` to `YES`.

Troubleshooting cloud storage configuration issues

The following sections may help you troubleshoot configuration issues.

See [“NetBackup Scalable Storage host properties unavailable”](#) on page 145.

See [“Cannot create a cloud storage”](#) on page 146.

See [“Data transfer to cloud storage server fails in the SSL mode”](#) on page 147.

See [“Amazon GovCloud cloud storage configuration fails in non-SSL mode”](#) on page 147.

See [“Data restore from the Google Nearline storage class may fail”](#) on page 147.

See [“Fetching storage regions fails with authentication version V2”](#) on page 148.

NetBackup Scalable Storage host properties unavailable

If the NetBackup CloudStore Service Container is not active, the **Scalable Storage** host properties are unavailable. Either of the following two symptoms may occur:

- The **Scalable Storage** properties for a media server are unavailable
- A pop-up box may appear that displays an **“Unable to fetch Scalable Storage settings”** message.

You should determine why the NetBackup CloudStore Service Container is inactive, resolve the problem, and then start the Service Container.

See [“NetBackup CloudStore Service Container startup and shutdown troubleshooting”](#) on page 154.

See [“Stopping and starting the NetBackup CloudStore Service Container”](#) on page 153.

Cannot create a cloud storage

If you cannot create a cloud storage in NetBackup, verify the following:

- The NetBackup `cacert.pem` file is present on both NetBackup primary and media server in following locations:
 - UNIX/Linux - `/usr/opensv/var/webtruststore`
 - Windows - `<install_path>/var/webtruststore`

On media server versions 7.7.x to 8.1.2, if the NetBackup `cacert.pem` file is not present, run the `nbcertcmd -getCACertificate` on the primary server. After running this command, restart the NetBackup CloudStore Service Container. See the *NetBackup Commands Reference Guide* for a complete description of the command.

Note: This NetBackup `cacert.pem` file is a NetBackup-specific file. This file includes the CA certificates generated by the NetBackup authorization service.

- The NetBackup `cacert.pem` file is same on the NetBackup primary and media server.
- For media server versions 7.7.x to 8.1.2, the machine certificate is present in following locations:
 - UNIX/Linux - `/usr/opensv/var/vxss/credentials`
 - Windows - `<install_path>/var/vxss/credentials`

If the security certificate is not present, run the `bpbaz -ProvisionCert` on the primary server. After running this command, restart the NetBackup CloudStore Service Container on the primary and media server.
 See [“Deploying host name-based certificates”](#) on page 87.
- For media server versions 7.7.x to 8.1.2, the NetBackup CloudStore Service Container is active.
 See [“Stopping and starting the NetBackup CloudStore Service Container”](#) on page 153.
- The **Enable insecure communication with 8.0 and earlier hosts** option on the NetBackup primary server is selected if the media server is of the version

8.0 or earlier. The option is available in the **NetBackup web UI** in the **Settings > Global security > Secure communication**.

- On the media server, if the certificate deployment security level is set to Very High, automatic certificate deployment is disabled. An authorization token must accompany every new certificate request. Therefore, you must create an authorization token before deploying the certificates.
 See the “Creating authorization tokens” topic in the *NetBackup Security and Encryption Guide* for detailed steps.

Data transfer to cloud storage server fails in the SSL mode

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

Amazon GovCloud cloud storage configuration fails in non-SSL mode

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

To enable the SSL mode again, run the `csconfig` command with `-us` parameter to set the value of SSL to '2'.

See the [NetBackup Commands Reference Guide](#) for a complete description about the commands.

Data restore from the Google Nearline storage class may fail

Data restore from the Google Nearline storage class may fail, if your `READ_BUFFER_SIZE` in NetBackup is set to a value that is greater than the allotted read throughput. Google allots the read throughput based on the total size of the data that you have stored in the Google Nearline storage class.

Note: The default `READ_BUFFER_SIZE` is 100 MB.

The NetBackup bptm logs show the following error after the data restore from Google Nearline fails:

```
HTTP status: 429, Retry type: RETRY_EXHAUSTED
```

Google provides 4 MB/s of read throughput per TB of data that you store in the Google Nearline storage class per location. You should change the `READ_BUFFER_SIZE` value in NetBackup to match it to the read throughput that Google allots.

For example, if the data that you have stored in the Google Nearline storage class is 5 TB, you should change the `READ_BUFFER_SIZE` value to match it to the allotted read throughput, which equals to 20 MB.

Refer to the Google guidelines, for more information:

<https://cloud.google.com/storage/docs/nearline?hl=en>

<https://cloud.google.com/storage/docs/nearline?hl=en>

See “[NetBackup cloud storage server connection properties](#)” on page 106.

Fetching storage regions fails with authentication version V2

When you use authentication version V2, if fetching storage regions step fails with pop-up error `Unable to process request (228)`, perform the following troubleshooting steps:

Ensure that `nbsl` and `nbwmc` services are up and running.

Enable `nbwmc` logs and in the `nblog.conf` file, increase verbosity to the highest level. Try fetching regions once again.

See “[NetBackup cloudstore.conf configuration file](#)” on page 84.

If the issue persists, look for cURL error in `csconfig` logs. The cURL error code helps you to find the root cause of the issue.

Some of the erroneous configuration scenarios can be:

- If the cURL error indicates that issue is caused due to invalid authentication URL, ensure that identity API version 2 endpoint (`v2.0/tokens`) is used for authentication.
 For example, `http://mycloud.xyz.com.com:5000/v2.0/tokens` must be used to authenticate instead of `https://mycloud.xyz.com:5000`.
- If the cURL error indicates that the issue is caused due to non-CA signed certificate, add a self-signed certificate to `cacert.pem` for *authentication* as well as *storage endpoint* (in case they are hosted separately).

Backup from snapshot parent jobs are failing with the status code 160

Check if the Snapshot Manager is accessible on port 443 from the media servers corresponding to the storage server that are configured in the protection plan.

Resolve the network issue by adding appropriate entries in the `/etc/hosts` file on the media server.

Troubleshooting cloud storage operational issues

The following sections may help you troubleshoot operational issues.

See [“NetBackup Scalable Storage host properties unavailable”](#) on page 145.

See [“Cloud storage backups fail”](#) on page 149.

See [“A restart of the nbcssc \(on legacy media servers\), nbwmc, and nbsl processes reverts all cloudstore.conf settings”](#) on page 154.

See [“NetBackup CloudStore Service Container startup and shutdown troubleshooting”](#) on page 154.

Cloud storage backups fail

See the following topics:

- [Accelerator backups fail](#)
- [Backups fail after the WRITE_BUFFER_SIZE is increased](#)
- [The storage volume was created by the cloud vendor interface](#)
- [The NetBackup CloudStore Service Container is not active](#)
- [Backups may fail if the Use any available media server option is selected](#)
- [Cloud backup and restore operations fail with error code 83 or error code 2106](#)
- [Cloud storage backup fails for certificate issues](#)
- [Backup jobs to Amazon S3 complaint cloud storage fail with status 41](#)

Accelerator backups fail

A message similar to the following is in the job details:

```
Critical bptm(pid=28291) accelerator verification failed: backupid= host_name_1373526632,
offset=3584, length=141976576, error= 2060022, error message: software error
Critical bptm(pid=28291) image write failed: error 2060022: software error
Error bptm(pid=28291) cannot write image to disk, Invalid argument end writing;
```

```
write time: 0:02:31
Info bptm(pid=28291) EXITING with status 84
Info bpbkar(pid=6044) done. status: 84: media write error media write error(84)
```

This error may occur in the environments that have more than one cloud storage server. It indicates that NetBackup Accelerator backups of a client to one cloud storage server were later directed to a different cloud storage server.

For Accelerator backups to cloud storage, ensure the following:

- Always back up each client to the same storage server. Do so even if the other storage server represents storage from the same cloud storage vendor.
- Always use the same backup policy to back up a client, and do not change the storage destination of that policy.

Backups fail after the WRITE_BUFFER_SIZE is increased

If the cloud storage server `WRITE_BUFFER_SIZE` property exceeds the total swap space of the computer, backups can fail with a status 84.

Adjust the `WRITE_BUFFER_SIZE` size to a value lower than the computer's total swap space to resolve this issue.

The storage volume was created by the cloud vendor interface

A message similar to the following is in the job details:

```
Info bptm(pid=xxx) start backup
Critical bptm(pid=xxxx) image open failed: error 2060029: authorization failure
Error bpbrm(pid=xxxx) from client gabby: ERR - Cannot write to STDOUT. Errno = 32:
Broken pipe
Info bptm(pid=xxxx) EXITING with status 84
```

A message similar to the following appears in the `bptm` log file:

```
Container container_name is not Cohesity container or tag data error,
fail to create image. Please make sure that the LSU is created by
means of NBU.
```

This error indicates that the volume was created by using the cloud storage vendor's interface.

Use the NetBackup web UI to create the volume on the cloud storage. NetBackup applies a required partner ID to the volume. If you use the vendor interface to create the container, the partner ID is not applied.

To resolve the problem, use the cloud storage vendor's interface to delete the container. In NetBackup, delete the disk pool and then recreate it with web UI.

See [“Viewing cloud storage job details”](#) on page 127.

See [“NetBackup cloud storage log files”](#) on page 142.

The NetBackup CloudStore Service Container is not active

This is applicable to media server versions 7.7.x to 8.1.2 only.

If the NetBackup CloudStore Service Container is not active, backups cannot be sent to the cloud storage.

NetBackup does not validate that the CloudStore Service Container is active when you use NetBackup commands to configure NetBackup cloud storage. Therefore, any backups that initiate in such a scenario fail.

See [“NetBackup CloudStore Service Container startup and shutdown troubleshooting”](#) on page 154.

Backups may fail if the Use any available media server option is selected

While you configure a cloud storage server, you must ensure that the media server and the primary server are of the same version.

Note: This limitation does not apply to the existing cloud storage servers.

Cloud backups may fail in the following scenario:

You selected **Use any available media server** while you configured the storage unit and NetBackup uses a media server with version different than the primary server version during cloud storage configuration.

To resolve this issue, do the following:

Select **Only use the following media servers** while you configure the storage unit and select the media server with a version same as primary server from the **Media Servers** pane.

See [“Troubleshooting cloud storage operational issues”](#) on page 149.

Cloud backup and restore operations fail with error code 83 or error code 2106

The cloud backups and restore operations failing with error code 83 or error code 2106 may occur due to any one of the following reasons:

- The media server's date and time settings are skewed (not in sync with the GMT/UTC time).
- The storage server credentials that are provided are incorrect.

Perform the following:

Change the media server's date and time settings so that it is in sync with the GMT/UTC time.

Update the storage server credentials. Use the `tpconfig` command to update the credentials. For more information, see the *NetBackup Commands Reference Guide*.

Cloud storage backup fails for certificate issues

If the cloud storage backups fails because of certificate issues, verify the following:

- The NetBackup `cacert.pem` file is present on both NetBackup primary and media server in following locations:

- UNIX/Linux - `/usr/opensv/var/webtruststore`
- Windows - `<install_path>/var/webtruststore`

For media server versions 7.7.x to 8.1.2, if the NetBackup `cacert.pem` file is not present, run the `nbcertcmd -getCACertificate` on the primary server. After running this command, restart the NetBackup CloudStore Service Container. See the *NetBackup Commands Reference Guide* for a complete description of the command.

Note: This NetBackup `cacert.pem` file is a NetBackup-specific file. This file includes the CA certificates generated by the NetBackup authorization service.

- The NetBackup `cacert.pem` file is same on the NetBackup primary and media server.
- For media server versions 7.7.x to 8.1.2, the machine certificate is present in following locations:

- UNIX/Linux - `/usr/opensv/var/vxss/credentials`
- Windows - `<install_path>/var/vxss/credentials`

If the security certificate is not present, run the `bnpbaz -ProvisionCert` on the primary server. After running this command, restart the NetBackup CloudStore Service Container on the primary and media server.

See [“Deploying host name-based certificates”](#) on page 87.

- For media server versions 7.7.x to 8.1.2, the NetBackup CloudStore Service Container is active.
See [“Stopping and starting the NetBackup CloudStore Service Container”](#) on page 153.
- The **Enable insecure communication with 8.0 and earlier hosts** option on the NetBackup primary server is selected if the media server is of the version

8.0 or earlier. The option is available at **Settings > Global Security > Secure Communication** tab.

- On the media server, if the certificate deployment security level is set to Very High, automatic certificate deployment is disabled. An authorization token must accompany every new certificate request. Therefore, you must create an authorization token before deploying the certificates.
 See the 'Creating authorization tokens' topic in the *NetBackup™ Security and Encryption Guide* for detailed steps.

Backup jobs to Amazon S3 complaint cloud storage fail with status 41

NetBackup consumes the available bandwidth to its maximum potential and pushes the requests accordingly, however the Amazon S3 complaint cloud is not able to process the number requests.

The cloud vendor returns error 503 to slow down the requests and the backup job fails with the following errors:

- In the media server `bptm` logs:

```
bptm:4940:<media_server_name>: AmzResiliency:
AmzResiliency::getRetryType cURL error: 0, multi cURL error: 0,
HTTP status: 503, XML response: SlowDown, RetryType:
RETRY_EXHAUSTED
```
- In the media server `bpbrm` logs:

```
bpbrm Exit: client backup EXIT STATUS 41: network connection timed
out
```

This issue arises only if higher bandwidth is available between NetBackup and the cloud storage.

To troubleshoot you can perform one of the following:

- Configure bandwidth throttling to reduce the number of requests.
 See [“NetBackup cloud storage server connection properties”](#) on page 106.
- Reduce the number of read/write buffers.
 See [“NetBackup cloud storage server bandwidth throttling properties”](#) on page 102.
- Talk to your cloud vendor to increase the number of parallel requests limit. This might incur extra cost.

Stopping and starting the NetBackup CloudStore Service Container

This is applicable to media server versions 7.7.x to 8.1.2 only.

Stop and start the NetBackup CloudStore Service Container (`nbcssc`) service.

See [“About the NetBackup CloudStore Service Container”](#) on page 81.

See [“NetBackup CloudStore Service Container startup and shutdown troubleshooting”](#) on page 154.

To start or stop the CloudStore Service Container

- 1** In the NetBackup web UI, click the Activity monitor.
- 2** Click the **Daemons** tab.
- 3** Locate the **nbcssc** service.
- 4** Click **Actions > Stop** or **Actions > Start**.

A restart of the nbcssc (on legacy media servers), nbwmc, and nbsl processes reverts all cloudstore.conf settings

Missing entries and comments are not allowed in the `cloudstore.conf` file. If you remove or comment out values in the `cloudstore.conf` file, a restart of the `nbcssc` (on older media servers), `nbwmc`, and `nbsl` processes on the media servers returns all settings to their default values.

NetBackup CloudStore Service Container startup and shutdown troubleshooting

This is applicable for media server versions 7.7.x to 8.1.2 only.

See the following topics:

- [Security certificate not provisioned](#)
- [Security mode changed while service is active](#)

Security certificate not provisioned

The NetBackup media servers that you use for cloud storage must have a security certificate provisioned. If not, the CloudStore Service Container cannot start. Verify that the certificate exists.

See [“NetBackup CloudStore Service Container security certificates”](#) on page 82.

NetBackup 7.7 to
8.1.2

If a certificate does not exist, create one from the NetBackup primary server.

See [“NetBackup CloudStore Service Container security certificates”](#) on page 82.

Security mode changed while service is active

Do not change the security mode of the NetBackup CloudStore Service Container while the service is active. If the security mode is changed while the service is active, you may encounter service startup or service shutdown problems. Be sure to stop the service in the same mode it was started.

See “[NetBackup CloudStore Service Container security modes](#)” on page 83.

See “[Stopping and starting the NetBackup CloudStore Service Container](#)” on page 153.

bptm process takes time to terminate after cancelling GLACIER restore job

During Amazon GLACIER restores on UNIX media servers, after canceling a restore job for images that are in GLACIER, the bptm process takes about 4 hours to terminate.

Workaround

You must manually kill the process.

Handling image cleanup failures for Amazon Glacier vault

The topic describes how to handle image cleanup failures for Amazon Glacier vault when vault lock policy is applied to the vault. Image cleanup fails when retention period set in the NetBackup policy is less than the period enforced by the vault lock policy applied on the Amazon Glacier vault storage unit.

To clean up image failures, see

<https://support.cohesity.com/s/article/article-100042245>.

Cleaning up orphaned archives manually

There may be instances where you cannot clean up orphaned images in Amazon Glacier vault due to the absence of a metadata object. A metadata object contains mapping information between data objects and NetBackup images.

To manually clean up orphaned archives in Amazon Glacier vault, see

<https://support.cohesity.com/s/article/article-100042314>.

Restoring from Amazon Glacier vault spans more than 24 hours for single fragment

Archives stored in Amazon Glacier vault, once retrieved, are available for download for only 24 hours. If your NetBackup restore job (for images residing in Amazon

Glacier vault) takes more than 24 hours to download a single fragment, the restore job may fail while reading the image. For example, if your fragment size is 512 GB and restore speed is less than 50 Mbps, the restore will fail.

To recover from this situation, do one of the following:

- Use a checkpoint restore.
- Start a restore for the remaining files.
- Duplicate the image with lesser fragment size.

Restoring from GLACIER_VAULT takes more than 24 hours for Oracle databases

Oracle forms a restore job, so that first the data files are restored (one job per data file) and then every set of archive logs (one restore job per set of logs) associated with the data files is restored. This causes the Oracle restore jobs to run five restore jobs in succession (when one restore job gets over, the next one automatically starts). Since every new restore job with data in a vault in Amazon Glacier cloud storage requires minimum four hours to retrieve the data to bring it on premise, this causes the Oracle data file restore jobs to run for 24 hours or longer.

There are two options to perform the recovery:

Using the NetBackup for Oracle recovery wizard

Increase the **Number of parallel streams for restore and recover** to the number of backup requests that are required. For example 10. You can set this number higher since Oracle RMAN will only use the required number of streams.

See section About NetBackup for Oracle restores in the *NetBackup for Oracle Administrator's Guide*.

Using the RMAN template

This procedure takes a longer time than the earlier mentioned method.

1. Determine the log sequence and thread numbers required for the recovery step (restoring the archive logs). This can be done by looking at Oracle or by looking at the backup jobs.
2. Create an RMAN script and allocate the required number of channels to perform the restore of the archive logs.

For example: Consider a “run” block where 8 channels are allocated and restored sequence numbers 1373 – 1380

```
RMAN> run
```

```
{ allocate channel ch00 type 'SBT_TAPE' PARMS
'SBT_LIBRARY=/bp/bin/libobk.so64';

allocate channel ch01 type 'SBT_TAPE' PARMS
'SBT_LIBRARY=/bp/bin/libobk.so64';

allocate channel ch02 type 'SBT_TAPE' PARMS
'SBT_LIBRARY=/bp/bin/libobk.so64';

allocate channel ch03 type 'SBT_TAPE' PARMS
'SBT_LIBRARY=/bp/bin/libobk.so64';

allocate channel ch04 type 'SBT_TAPE' PARMS
'SBT_LIBRARY=/bp/bin/libobk.so64';

allocate channel ch05 type 'SBT_TAPE' PARMS
'SBT_LIBRARY=/bp/bin/libobk.so64';

allocate channel ch06 type 'SBT_TAPE' PARMS
'SBT_LIBRARY=/bp/bin/libobk.so64';

allocate channel ch07 type 'SBT_TAPE' PARMS
'SBT_LIBRARY=/bp/bin/libobk.so64';
```

Restore the archive log from sequence 1373 thread 1 until sequence 1380 thread 1;

```
release channel ch00;

release channel ch01;

release channel ch02;

release channel ch03;

release channel ch04;

release channel ch05;

release channel ch06;

release channel ch07;

}
```

3. Using the NetBackup for Oracle client, start NetBackup Backup, Archive, and Restore interface or create another script to restore the data file or files. If you're restoring more than one data file, you may need to increase the number of streams if each data file is in a different image.
4. Start the restore of the data files and archive logs to run in parallel.

5. Perform the recovery of the database or data files using the NetBackup Backup, Archive, and Restore interface or by using another script.

See the *NetBackup for Oracle Administrator's Guide*.

Troubleshooting failures due to missing Amazon IAM permissions

If the AWS credential provided in NetBackup cloud configuration does not have S3 or Glacier related permission, you could see failures or errors at various stages of configuration, backup, and restore.

Some error messages are clearly described and identifiable in the NetBackup web UI, while others are vague.

Amazon displays the `AccessDeniedException` error message. To decipher this error message, you need to check the log files to check for the missing permission.

- List Vault or List Bucket permission (`glacier:ListVaults`) missing.

This error occurs while creating a storage server. If you are using the CLI, `tpcommand` to add credential fails.

Check the `tpcommand` logs for `AccessDeniedException`, for example,

```
amazon: Json:
{"code":"AccessDeniedException","type":"Client","message":"User:
arn:aws:iam::326221795898:user/ReadOnly_user is not authorized to
perform: glacier:ListVaults on resource:
arn:aws:glacier:ap-south-1:326221795898:vaults/"} 16:17:52.139
[7388.4424] <2> magmavml.abc.xyz.qwe.com: AmzVaultApi:
json_string({"code":"AccessDeniedException","type":"Client","message":"User:
arn:aws:iam::326221795898:user/ReadOnly_user is not authorized to
perform: glacier:ListVaults on resource:
arn:aws:glacier:ap-south-1:326221795898:vaults/"}) 16:17:52.139
[7388.4424] <16> magmavml.abc.xyz.qwe.com:
```

- Create Vault or Create Bucket permission (`glacier:CreateVault` or `glacier:DescribeVault`) missing.

This error occurs while creating a disk pool using the NetBackup Administrator console. If you are using the CLI, `nbdevconfig` command fails.

Check the `nbrrms` log for `AccessDeniedException`, for example,

```
amazon_raw:: AmzVaultApi: Error: server error code
AccessDeniedException, User:
arn:aws:iam::326221795898:user/ReadOnly_user is not authorized to
perform: glacier:CreateVault on resource:
arn:aws:glacier:ap-south-1:326221795898:vaults/fail-to-create,
httpcode [403] returning [2060037],11:STS Service,1Post Archive
```

or S3 Object permission missing - backup will fail in activity monitor.

- Upload archives permission (`glacier:UploadArchive`) missing.
This error occurs while backing up archives. The backup jobs fail with permission error.

Check the `bptm` log for details, for example,

```
"code": "AccessDeniedException", "type": "Client", "message": "User:
arn:aws:iam::3234415151:user/XYZ is not authorized to perform:
glacier:UploadArchive on resource: LSTR-gtwy-00076(debug)."
```

- Retrieve job after archive permission (`glacier:InitiateJob`) missing.
This errors occurs after you start a restore.

Check the `bptm` log for details, for example,

```
"code": "AccessDeniedException", "type": "Client", "message": "User:
arn:aws:iam::3234415151:user/XYZ is not authorized to perform:
glacier:InitiateJob on resource: LSTR-gtwy-00076(debug)."
```

- Retrieve Archive or retrieve Object permission missing (`glacier:GetJobOutput`) missing.

This missing permission causes the restore job to be in incomplete state if NetBackup cannot download archives after posting jobs.

Check the `bptm` log for details, for example,

```
"code": "AccessDeniedException", "type": "Client", "message": "User:
arn:aws:iam::3234415151:user/XYZ is not authorized to perform:
glacier:GetJobOutput on resource: LSTR-gtwy-00076(debug)."
```

- Delete Archive or Delete Object permission (`glacier>DeleteArchive`) missing.
This missing permission causes the image cleanup or image expiry process to fail.

Check the `bptm` log for details, for example,

```
"code": "AccessDeniedException", "type": "Client", "message": "User:
arn:aws:iam::3234415151:user/XYZ is not authorized to perform:
glacier>DeleteArchive on resource: LSTR-gtwy-00076(debug)."
```

Restore job fails if the restore job start time overlaps with the backup job end time

If you trigger a restore job within a few seconds of the backup job completion, the restore job fails with the following error:

Standard policy restore error

The restore job in such scenario fails because the cloud provider requires time to update the parameters required for performing a restore. Thus, trigger the restore after a few minutes from the backup job completion.

Post processing fails for restore from Azure archive

When post processing for restore from Azure archive fails, blobs are not moved from the Hot tier to Archive tier post restore.

To move the blobs from Hot tier to Archive tier follow the steps:

- Use the list blob operation and get a list blobs with prefix REHYDRATE_PENDING. The blob names are returned in format - REHYDRATE_PENDING/<image_name>
- Search for blobs with <image_name>/ as prefix and filter with the blob names in integer format after the prefix.

For example :

Consider image name as imagename_1544519515_C1_F1

Blob to pick for post processing - imagename_1544519515_C1_F1/21

Blob not to be picked up -

imagename_1544519515_C1_F1/imagename_1544519515/0

- Use the set blob tier operation on blob to change the access tier of the blobs returned from above step from hot access tier to archive access tier.

Note: Do not move the `META_BLOCK_MAP_FILE` and `META_IMAGE_PROPERTIES` and blobs to the archive tier.

- After you successfully, move the blob to archive access tier, delete the blob with prefix REHYDRATE_PENDING using the delete blob operation.

Troubleshooting Amazon Snowball and Amazon Snowball Edge issues

Disk pool creation fails

Disk pool creation fails when the cloud storage properties are changed to Amazon Snowball endpoint. The following error is encountered:

```
No Volumes found.
```

To troubleshoot:

Ensure that the `OFFLINE_TRANSFER_MODE` storage server property is set to `PROVIDER_API`.

Restore fails

Restore fails with the following error:

```
The specified key does not exist.
```

The image to be restored was not successfully imported to cloud. Re-run the duplication-to-cloud operation for that image and perform the restore.

Run the `bpduplicate` command. See the [NetBackup Command Reference Guide](#).

Import to cloud fails

Run the duplication-to-cloud operation for that image. Use the `bpduplicate` command. See the [NetBackup Command Reference Guide](#).

For any other issues, ensure that the configuration is done properly. Refer to the [NetBackup with Amazon Snowball and Snowball Edge Configuration Checks](#) tech note.

Index

A

- Add at least one index marker 71
- Amazon
 - glacier vault 39–40
- amazon
 - virtual private cloud 29
- amazon (S3)
 - permissions 18
- Amazon Amazon S3 Intelligent Tiering
 - restore 45
- Amazon GLACIER
 - long-term retention 32
- Amazon Glacier 31
- Amazon Glacier Deep Archive 31
- Amazon Glacier Vault 31
- Amazon IAM roles 45
- Amazon S3
 - about 15
 - credential broker details 26
 - requirements 16
- Amazon S3 Intelligent Tiering
 - backup 44
 - cloud tiering 42
- Amazon Snowball 50
 - configuring with Amazon S3 API interface 53
 - configuring with Amazon Snowball client 52
- Amazon Snowball Edge 50
 - configuring with file interface 56
 - configuring with S3 API interface 58

B

- backups fail
 - The NetBackup CloudStore Service Container is not active 151
 - Use any available media server option 151
- bandwidth
 - throttling 102
- bpstinfo command
 - operational notes 131

C

- Certificate Authority (CA) 89
- cloud
 - storage unit properties 118
- cloud primary host 98
- cloud storage
 - Amazon S3 API type 15
 - configuring 72
 - Microsoft Azure API type 62
- cloud storage server
 - about 94
 - bandwidth properties 102
 - encryption properties 112
- CloudStore Service Container
 - security mode changed while service is active 155
 - security modes 83
 - startup and shutdown troubleshooting 154
- configuration
 - disk pool configuration wizard 114
- configuring a deduplication storage unit 117
- configuring cloud storage 72

D

- Deduplication storage unit
 - Only use the following media servers 118
 - Use any available media server 118
- Disk type 118
- Dynamic Host Configuration Protocol (DHCP) 88

E

- encryption
 - properties 112
- external KMS 93

F

- Features and functionality 8
- FlashBackup policy
 - Maximum fragment size (storage unit setting) 119

G

glacier vault
 back up 39
 restore 40

H

host ID-based certificates
 deploying with a token 89
 deploying without a token 89
host name-based certificates
 deploying 88
hotfix 88

I

IAM Role 48

J

job ID search in unified logs 140

M

Maximum concurrent jobs 119
Maximum fragment size 119
Microsoft Azure
 about 62
 configuration options 64
 configuration options (advanced) 66
 requirements 63

N

NetBackup
 hotfix 88
NetBackup Accelerator
 about 120
NetBackup Scalable Storage host properties
 unavailable 145
NetBackup Service Layer (NBSL) 88

O

Optimized Synthetic backups
 about 120

P

Preferences
 encryption 113
 throttling 112

private clouds

 Amazon S3-compatible cloud providers 27

properties

 bandwidth 102

 encryption 112

R

requirements 73

S

Scalable Storage host properties unavailable 145

storage server. *See* cloud storage server

storage unit

 configuring for deduplication 117

 properties for cloud 118

Storage unit name 118

Storage unit type 118

U

unified logging 136

 format of files 137

V

virtual private cloud 29

VPC 29

vxlogview command 137

 with job ID option 140