

Veritas NetBackup™ Cloud Administrator's Guide

UNIX, Windows, Linux

Release 8.1.2

VERITAS™

Veritas NetBackup™ Cloud Administrator's Guide

Last updated: 2019-03-20

Legal Notice

Copyright © 2019 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies LLC and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. VERITAS TECHNOLOGIES LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

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

<https://www.veritas.com/support>

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

<https://sort.veritas.com/documents>

Documentation feedback

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

NB.docs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	About NetBackup cloud storage	8
	About cloud storage features and functionality	8
	About the catalog backup of cloud configuration files	12
	About support limitations for NetBackup cloud storage	12
Chapter 2	About the 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
	Amazon S3 cloud storage provider options	17
	Amazon S3 cloud storage options	23
	Amazon S3 advanced server configuration options	25
	Amazon S3 credentials broker details	28
	About private clouds from Amazon S3-compatible cloud providers	30
	About Amazon S3 storage classes	31
	Amazon virtual private cloud support with NetBackup	32
	About protecting data in Amazon for long-term retention	33
	Protecting data using Amazon's cloud tiering	45
	About NetBackup character restrictions for Amazon S3 cloud connector	48
	About Microsoft Azure cloud storage API type	50
	Microsoft Azure cloud storage vendors certified for NetBackup	50
	Microsoft Azure storage type requirements	50
	Microsoft Azure cloud storage provider options	51
	Microsoft Azure advanced server configuration options	55
	About OpenStack Swift cloud storage API type	57
	OpenStack Swift cloud storage vendors certified for NetBackup	58
	OpenStack Swift storage type requirements	58
	OpenStack Swift cloud storage provider options	59
	OpenStack Swift storage region options	63
	OpenStack Swift add cloud storage configuration options	65

	OpenStack Swift proxy settings	65
Chapter 3	Configuring cloud storage in NetBackup	67
	Before you begin to configure cloud storage in NetBackup	68
	Configuring cloud storage in NetBackup	69
	Cloud installation requirements	71
	Scalable Storage properties	71
	Configuring advanced bandwidth throttling settings	73
	Advanced bandwidth throttling settings	74
	Cloud Storage properties	76
	Adding a cloud storage instance	78
	Changing cloud storage host properties	79
	Deleting a cloud storage host instance	80
	About the NetBackup CloudStore Service Container	81
	NetBackup CloudStore Service Container security certificates	82
	NetBackup CloudStore Service Container security modes	84
	NetBackup cloudstore.conf configuration file	84
	Deploying host name-based certificates	88
	Deploying host ID-based certificates	90
	About data compression for cloud backups	91
	About data encryption for cloud storage	92
	About key management for encryption of NetBackup cloud storage	93
	About cloud storage servers	94
	About object size for cloud storage	95
	About the NetBackup media servers for cloud storage	97
	Using media server as NetBackup Cloud master host	98
	Configuring a storage server for cloud storage	100
	KMS database encryption settings	104
	Assigning a storage class to Amazon cloud storage	105
	Changing cloud storage server properties	106
	NetBackup cloud storage server properties	108
	NetBackup cloud storage server bandwidth throttling properties	109
	NetBackup cloud storage server connection properties	112
	NetBackup CloudCatalyst storage server properties	118
	NetBackup cloud storage server encryption properties	119
	About cloud storage disk pools	120
	Configuring a disk pool for cloud storage	120
	Saving a record of the KMS key names for NetBackup cloud storage encryption	129

	Adding backup media servers to your cloud environment	131
	Configuring a storage unit for cloud storage	132
	Cloud storage unit properties	133
	Configure a favorable client-to-server ratio	135
	Control backup traffic to the media servers	136
	About NetBackup Accelerator and NetBackup Optimized Synthetic backups	136
	Enabling NetBackup Accelerator with cloud storage	137
	Enabling optimized synthetic backups with cloud storage	138
	Creating a backup policy	140
	Changing cloud storage disk pool properties	141
	Cloud storage disk pool properties	142
	Managing Certification Authorities (CA) for NetBackup Cloud	144
Chapter 4	Monitoring and Reporting	148
	About monitoring and reporting for cloud backups	148
	Viewing cloud storage job details	150
	Viewing the compression ratio	150
	Viewing NetBackup cloud storage disk reports	151
	Displaying KMS key information for cloud storage encryption	152
Chapter 5	Operational notes	155
	NetBackup bpstsinfo command operational notes	155
	Unable to configure additional media servers	156
	Cloud configuration may fail if NetBackup Access Control is enabled	156
	Deleting cloud storage server artifacts	157
Chapter 6	Troubleshooting	158
	About unified logging	158
	About using the vxlogview command to view unified logs	159
	Examples of using vxlogview to view unified logs	160
	About legacy logging	161
	Creating NetBackup log file directories for cloud storage	163
	NetBackup cloud storage log files	163
	Enable libcurl logging	166
	NetBackup Administration Console fails to open	167
	Troubleshooting cloud storage configuration issues	168
	NetBackup Scalable Storage host properties unavailable	168
	Connection to the NetBackup CloudStore Service Container fails	169

Cannot create a cloud storage disk pool	170
Cannot create a cloud storage	171
Data transfer to cloud storage server fails in the SSL mode	172
Amazon GovCloud cloud storage configuration fails in non-SSL mode	172
Data restore from the Google Nearline storage class may fail	173
Backups may fail for cloud storage configurations with Frankfurt region	173
Backups may fail for cloud storage configurations with the cloud compression option	174
Fetching storage regions fails with authentication version V2	174
Troubleshooting cloud storage operational issues	175
Cloud storage backups fail	175
Stopping and starting the NetBackup CloudStore Service Container	179
A restart of the nbcssc (on legacy media servers), nbwmc, and nbsl processes reverts all cloudstore.conf settings	180
NetBackup CloudStore Service Container startup and shutdown troubleshooting	180
bptm process takes time to terminate after cancelling GLACIER restore job	181
Handling image cleanup failures for Amazon Glacier vault	181
Cleaning up orphaned archives manually	181
Restoring from Amazon Glacier vault spans more than 24 hours for single fragment	182
Restoring from GLACIER_VAULT takes more than 24 hours for Oracle databases	182
Troubleshooting failures due to missing Amazon IAM permissions	184
Restore job fails if the restore job start time overlaps with the backup job end time	189
Index	190

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 Veritas OpenStorage.

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

Table 1-1 Features and functionality

Feature	Details
Configuration Wizard	A Cloud Storage Server Configuration wizard is incorporated to facilitate the cloud storage setup and storage provisioning. Cloud storage provisioning now happens entirely through the NetBackup interface.
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 LZOP (with compression level 3).

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

Feature	Details
Encryption	<p>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.</p> <p>The encryption feature uses an AES 256 cipher feedback (CFB) mode encryption.</p>
Throttling	<p>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.</p> <p>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.</p> <p>NetBackup Cloud Storage Throttling lets you configure and control the following:</p> <ul style="list-style-type: none"> ■ 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	<p>The NetBackup Cloud Storage metering reports enable you to monitor data transfers within NetBackup Cloud Storage.</p> <p>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.</p> <p>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.</p> <p>Metering reports are generated through NetBackup OpsCenter.</p>

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

Feature	Details
Cloud Storage service	<p>This is applicable to media server versions 7.7.x to 8.1.2 only.</p> <p>The NetBackup CloudStore Service Container (<i>nbcssc</i>) process performs the following functions:</p> <ul style="list-style-type: none"> ■ Generates the metering information for the metering plug-in ■ Controls the network bandwidth usage with the help of the throttling plug-in <p>Note: For NetBackup media server versions beyond 8.1.2, these Cloud Storage functions are performed by the NetBackup Service Layer (<i>nbsl</i>) service.</p> <p>On Windows, it is a standard service installed by NetBackup. On UNIX, it runs as a standard daemon.</p> <p>The NetBackup CloudStore Service Container (<i>nbcssc</i>) uses certificate-based authentication. The authentication method used in previous releases (legacy authentication) is disabled by default. Veritas recommends that you upgrade media servers configured as a cloud storage server to NetBackup 8.1 or later.</p> <p>If you cannot upgrade these servers, use the Enable insecure communication with 8.0 and earlier hosts option on the NetBackup master server. The option is available in the NetBackup Administration Console on the Security Management > Global Security Settings > Secure Communication tab.</p>
NetBackup Web Management Console	<p>The NetBackup Web Management Console (<i>nbwmc</i>) process manages requests for certificate and host management.</p> <p>This process now also controls the configuration parameters that are related to NetBackup Cloud Storage.</p> <p>The process is installed as a NetBackup service on Windows and runs as a standard daemon on UNIX.</p>

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

Feature	Details
NetBackup Service Layer	<p>The NetBackup Service Layer (<i>nbsl</i>) service facilitates the communication between NetBackup graphical user interface (UI) and the NetBackup logic. This service is required to run NetBackup OpsCenter which manages and monitors multiple NetBackup environments.</p> <p>This service is also required for Cloud Storage and now performs the following functions:</p> <ul style="list-style-type: none"> ■ Generates the metering information for the metering plug-in ■ Controls the network bandwidth usage with the help of the throttling plug-in <p>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>).</p>
Storage providers	<p>Veritas currently supports several cloud storage providers. More information is available about each of these vendors.</p> <p>See “About the cloud storage vendors for NetBackup” on page 14.</p>
OpsCenter Reporting	<p>Monitoring and reporting of the data that is sent to cloud storage is available through new cloud reports in OpsCenter. The cloud reports include:</p> <ul style="list-style-type: none"> ■ Job Success Rate: Success rate by backup job level across domains, clients, policies, and business level views filtered on cloud-based storage. ■ Data Expiring In Future: Data that expires each day for the next 7 days filtered on cloud-based storage. ■ Cloud Metering: Historical view of the data that is written to cloud per cloud provider. ■ Average Data Transfer Rate: Historical view of average data transfer rate to cloud per cloud provider. ■ Cloud Metering Chargeback: Ranking, forecast, and distribution view of the cost that is incurred on cloud-based storage per cloud provider. <p>Note: Among all Amazon S3-compatible cloud providers that NetBackup supports, OpsCenter supports monitoring and reporting of Amazon S3 and Amazon GovCloud (S3) only.</p> <p>Note: Where Amazon is the cloud service provider, OpsCenter cannot report on the data that MSDP cloud storage servers upload to the cloud.</p>

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`
- `CloudProvider.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 location:

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

UNIX `/usr/opensv/var/global/wmc/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 certificates of NetBackup supported Certificate Authorities (CA).

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 master 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:
<http://www.netbackup.com/compatibility>
<http://www.netbackup.com/compatibility>
- 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 Cloud Storage Configuration wizard, 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 the 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 OpenStack Swift 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:

http://jayne.dab.veritas.com/cl/hd/nbu/80/out/nbu_80_hd.html#cloud_storage_solutions

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

Vendors achieve certification by participating in the Veritas technology partners program. NetBackup can send backups to the storage that these vendors provide. Veritas 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 master compatibility list landing page:

<http://www.netbackup.com/compatibility>

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 50.
OpenStack Swift	See “About OpenStack Swift cloud storage API type” on page 57.

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 17.
Service host and endpoint configuration options	See “Amazon S3 cloud storage options” on page 23.
SSL, proxy, and HTTP header options	See “Amazon S3 advanced server configuration options” on page 25.
Credential broker options	See “Amazon S3 credentials broker details” on page 28.
Storage classes	See “About Amazon S3 storage classes” on page 31.

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 30.

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 using the Amazon S3 storage API as of the NetBackup 8.1.2 release:

http://jayne.clab.veritas.com/cl/hcl/nbu/80/out/nbu_80_hcl.html#cloud_storage_solutions

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

See [“Amazon S3 cloud storage provider options”](#) on page 17.

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

http://jayne.clab.veritas.com/cl/hcl/nbu/80/out/nbu_80_hcl.html#cloud_storage_solutions

http://jayne.clab.veritas.com/cl/hcl/nbu/80/out/nbu_80_hcl.html#cloud_storage_solutions

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

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: <ul style="list-style-type: none"> ■ 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.

Table 2-3 Amazon cloud storage requirements (*continued*)

Requirement	Details
Bucket names	<p>Veritas recommends 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.</p> <p>The following are the NetBackup requirements for bucket names in the US Standard region.</p> <ul style="list-style-type: none"> ■ 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): <ul style="list-style-type: none"> Period (.), underscore (_), and dash (-). Dash - <p><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 112.</p> <p>Note: The buckets are not available for use in NetBackup in the following scenarios: a) If you have created the buckets in a region that NetBackup does not support. b) The bucket name does not comply with the bucket naming convention.</p>
Number of disk pools	<p>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.</p>

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

Amazon S3 cloud storage provider options

[Figure 2-2](#) shows the **Cloud Storage Configuration Wizard** panel for Amazon S3 cloud storage.

Figure 2-1 Cloud Storage Server Configuration Wizard panel for Amazon

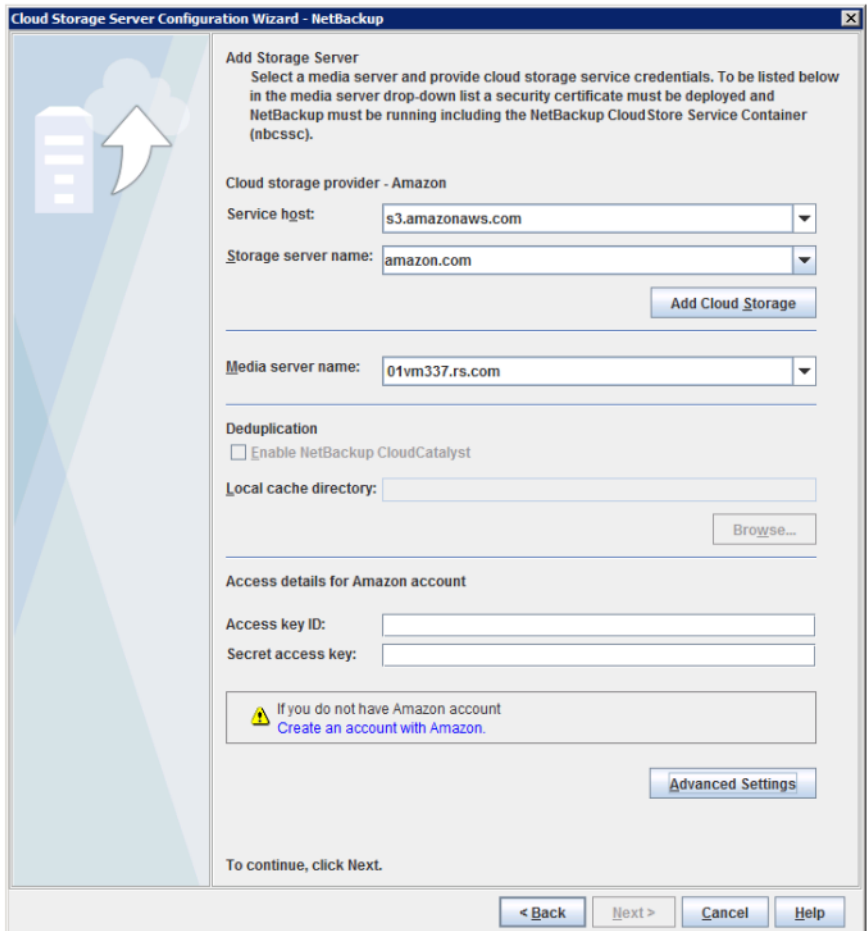


Figure 2-2 Cloud Storage Server Configuration Wizard panel for Amazon




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	<p>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.</p> <p>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.</p> <p>Note: Veritas recommends 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.</p> <p>Note: The Add Cloud Storage option is disabled for public clouds. You must use existing cloud storage.</p>
Add Cloud Storage	<p>To configure cloud deployment details, click Add Cloud 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.</p> <p>See “Amazon S3 cloud storage options” on page 23.</p> <p>Once the cloud storage is added, you cannot modify or delete it using the NetBackup Administration Console. However, you can modify or delete a storage server by using the <code>csconfig</code> command.</p> <p>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.</p> <p>See the <i>NetBackup Commands Reference Guide</i> for a complete description about these commands. The guide is available through the following URL:</p> <p>http://www.veritas.com/docs/DOC5332</p>

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

Field name	Required content
Media server name	<p>Select a NetBackup media server from the drop-down list. The drop-down list displays only NetBackup 8.1.2 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:</p> <p>See “About the NetBackup media servers for cloud storage” on page 97.</p> <p>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.</p> <p>To support cloud storage, a media server must conform to the following items:</p> <ul style="list-style-type: none"> ■ 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: http://www.netbackup.com/compatibility ■ The NetBackup Service Layer (<i>nbsl</i>) service must be running on all the media servers. The NetBackup Web Management Console (<i>nbwmc</i>) must be running on the master server. ■ For Amazon S3-compatible cloud providers, the media server must run a NetBackup 8.1.2 or later release. ■ The NetBackup media servers that you use for cloud storage must be the same NetBackup version as the master server.
Enter Credentials	<p><i>Applies to: Amazon GovCloud only.</i></p> <p>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.</p>
Use Credentials Broker	<p><i>Applies to: Amazon GovCloud only.</i></p> <p>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.</p>

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

Field name	Required content
Deduplication	<p>Enabling this option creates a CloudCatalyst storage server that can be used to upload deduplicated data to the cloud.</p> <p>This option is grayed out if any of the following cases are true:</p> <ul style="list-style-type: none"> ■ The selected media server does not have NetBackup 8.1 or later installed. ■ CloudCatalyst does not support the media server operating system. ■ CloudCatalyst does not support the cloud vendor. <p>See the NetBackup compatibility lists for support information: http://www.netbackup.com/compatibility</p> <p>For information about CloudCatalyst, see the <i>NetBackup Deduplication Guide</i>: http://www.veritas.com/docs/DOC5332</p>
Local cache directory	<p>Enter the mount path to be used as the storage path on the CloudCatalyst storage server.</p> <p>For example: <code>/space/mnt/esfs</code></p> <p>The deduplicated data is written to this local cache directory before it is uploaded to the cloud. The larger the cache, the more likely that NetBackup can service requests locally, avoiding cloud access to read and write.</p> <p>Notes:</p> <ul style="list-style-type: none"> ■ This path should be to a file system which is dedicated for CloudCatalyst cache use. Inaccurate cache eviction occurs if the path shares any storage with other data or applications. ■ NetBackup manages the files in the local cache directory. Users should not manually delete files in this directory.
Access key ID	<p><i>Does not apply for Amazon GovCloud if you select Use Credentials Broker.</i></p> <p>Enter the access key ID for your vendor account.</p> <p>If you do not have an account, click Create an account with the service provider link.</p>
Secret access key	<p><i>Does not apply for Amazon GovCloud if you select Use Credentials Broker.</i></p> <p>Enter the secret access key for your vendor account. It must be 100 or fewer characters.</p>

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

Field name	Required content
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 advanced server configuration options” on page 25.

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

Amazon S3 cloud storage options

The **Add Cloud Storage** dialog box appears when you click **Add Cloud Storage** on the wizard panel for Amazon S3 providers. It contains the following tabs:

General Settings tab See [Table 2-5](#) on page 23.

Region Settings tab See [Table 2-6](#) on page 25.

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 32.

Table 2-5 General Settings tab options

Option	Description
Provider type	The cloud storage provider. The following describes the state of this field: <ul style="list-style-type: none"> 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 Cloud Storage Server Configuration Wizard or change settings from the Cloud Storage host properties. It shows the host that you selected in the wizard or Cloud Storage host properties.

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

Option	Description
Service host	<p>Enter the cloud service provider host name.</p> <p>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.</p> <p>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'</p> <p>For custom instance, to use IPv6 endpoint, you must update or create a new instance with the IPv6 equivalent service host.</p> <p>Note: Do not prefix the service host name with 'http' or 'https'.</p> <p>Note: For VPC in default (US East (N. Virginia)) AWS region, use external-1.amazonaws.com as the service host.</p>
Service endpoint	<p>Enter the cloud service provider endpoint.</p> <p>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.</p> <p>You can leave it blank, if the cloud provider service can be accessed directly from the 'service.my-cloud.com' URL.</p>
HTTP port	<p>Enter the HTTP port with which you can access the cloud provider service in a non-secure mode.</p>
HTTPS port	<p>Enter the HTTPS port with which you can access the cloud provider service in a secure mode.</p>
Storage server name	<p>Enter a logical name for the cloud storage that you want to configure and access using NetBackup.</p> <p>Note: You can configure multiple storage servers that are associated with the same public or private cloud storage instance.</p>
Endpoint access style	<p>Select the endpoint access style for the cloud service provider.</p> <p>Path Style is the default endpoint access style.</p> <p>If your cloud service provider additionally supports virtual hosting of URLs, select Virtual Hosted Style.</p>

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 advanced server configuration options

The following tables describes the SSL, HTTP header configuration, and proxy server options that are specific to all Amazon S3-compatible cloud providers. These options appear on the **Advanced Server Configuration** dialog box.

Table 2-7 General Settings tab options

Option	Description
Use SSL	<p>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.</p> <ul style="list-style-type: none"> ■ 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. <p>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.</p> <p>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.</p> <p>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.</p>

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

Option	Description
HTTP Headers	<p>Specify appropriate value for the selected HTTP header. Click the Value column to see the drop-down list and select the value.</p> <ul style="list-style-type: none"> ■ 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 master key that it regularly rotates. <p>Note: If you have already enabled the encryption option while creating Amazon S3 cloud storage server, you do not need to enable this option. Because, the data is already encrypted before NetBackup sends it over the network.</p> ■ Storage class is configured at the time of creating the storage server. Once configured, storage class is non-editable.

Table 2-8 Proxy Settings tab options

Option	Description
Use Proxy Server	<p>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:</p> <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ HTTP Note: You need to provide the proxy credentials for HTTP proxy type. ■ SOCKS ■ SOCKS4 ■ SOCKS5 ■ SOCKS4A

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

Option	Description
Use Proxy Tunneling	<p>You can enable proxy tunneling for HTTP proxy type.</p> <p>After you enable Use Proxy Tunneling, HTTP CONNECT requests are sent from the cloud media server to the HTTP proxy server and the TCP connection is directly forwarded to the cloud back-end storage.</p> <p>The data passes through the proxy server without reading the headers or data from the connection.</p>
Authentication Type	<p>You can select one of the following authentication types if you are using HTTP proxy type.</p> <ul style="list-style-type: none"> ■ None— Authentication is not enabled. Username and password is not required. ■ NTLM—Username and password needed. ■ Basic—Username and password needed. <p>Username is the username of the proxy server</p> <p>Password can be empty. You can use maximum 256 characters.</p>

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

Amazon S3 credentials broker details

[Figure 2-3](#) shows the **Cloud Storage Configuration Wizard** credentials broker panel for Amazon GovCloud cloud storage. You add the credentials broker details when you configure a cloud storage server in NetBackup.

See [“Configuring a storage server for cloud storage”](#) on page 100.

The credentials broker details also appear in a **Cloud Storage Server Configuration** dialog box in which you can change the details.

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

Figure 2-3 Cloud Storage Server Configuration Wizard panel for Amazon

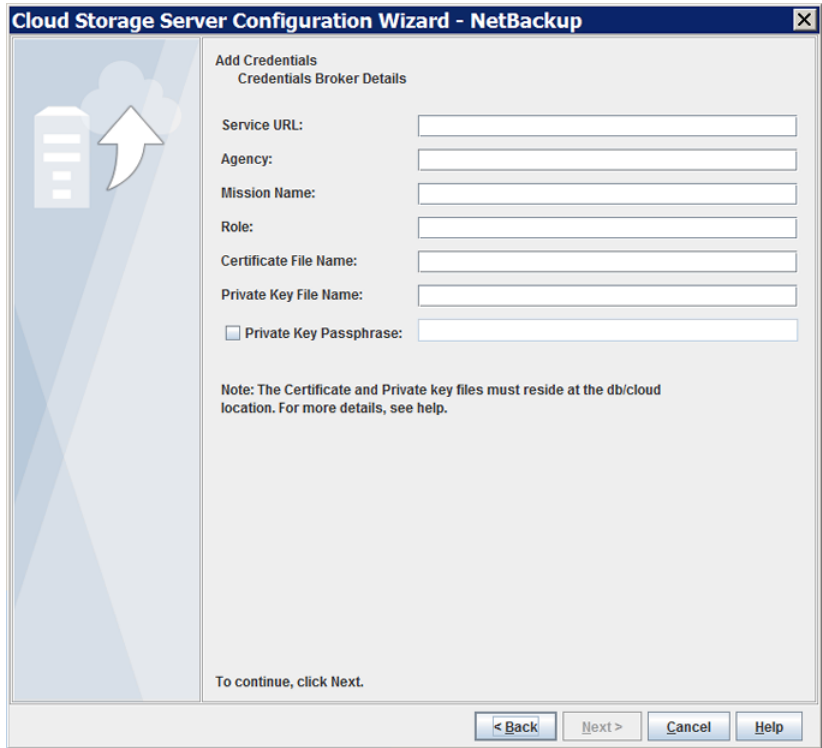


Table 2-9 describes the credential broker options for Amazon GovCloud.

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.
Role	Enter the role.
Certificate File Name	Enter the certificate file name.
Private Key File Name	Enter the private key file name.

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

Field	Description
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 Veritas 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

On the select media server panel of the **Cloud Storage Configuration Wizard**, click the **Advanced Settings** option. Then, in the **Advanced Server Configuration** dialog box, select the relevant options from the following: **Use SSL**, **Use Proxy Server**, **HTTP Headers**, and so on.

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 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 lifecycle) (IA stands for Infrequent Access.)**
- **GLACIER**
See [“About protecting data in Amazon Glacier”](#) on page 35.
- **GLACIER_VAULT**
See [“About protecting data in Amazon Glacier vault”](#) on page 39.
- **LIFECYCLE**
See [“Protecting data using Amazon's cloud tiering ”](#) on page 45.

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

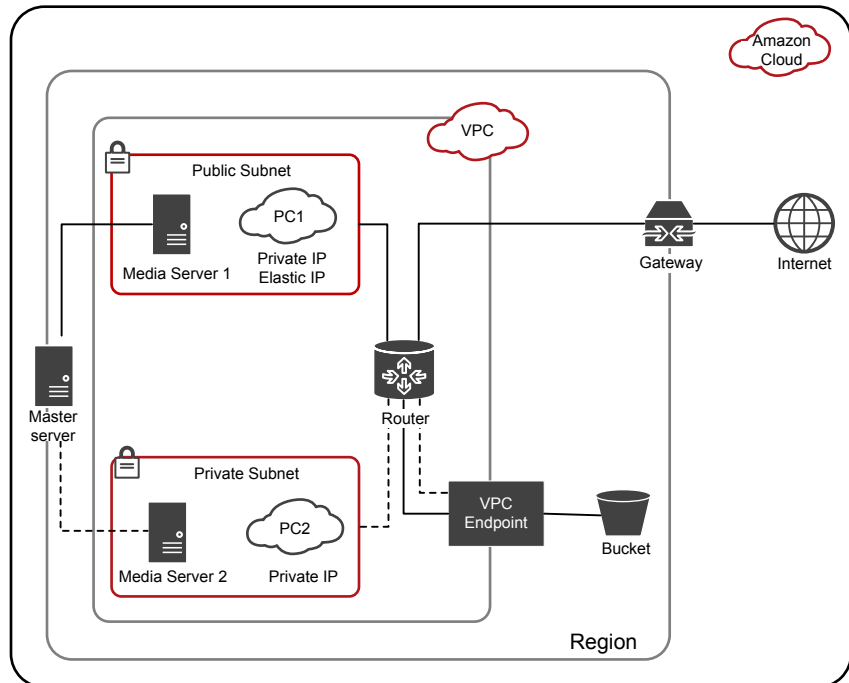
- If you do not select a specific storage class while you configure the Amazon S3 cloud storage
- If the backups were configured in an earlier NetBackup version

See [“Assigning a storage class to Amazon cloud storage”](#) on page 105.

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 master server locally or in the VPC environment. Ensure that the master 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. See [“Amazon S3 cloud storage options”](#) on page 23.
- 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 35.
- See [“About protecting data in Amazon Glacier vault”](#) on page 39.

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

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

GLACIER storage class

GLACIER storage class corresponds to uploading data through S3 end point and transitioning the data to Glacier storage using Amazon's zero day lifecycle policy that NetBackup creates.

For GLACIER storage class, the metadata is stored in STANDARD storage class.

GLACIER_VAULT storage class

GLACIER_VAULT storage class corresponds to uploading data using Amazon Glacier services to vault.

For GLACIER_VAULT storage class, the metadata is stored in STANDARD and GLACIER_VAULT storage classes.

GLACIER storage class

Cost of operation for GLACIER is approximately 2% higher than GLACIER_VAULT.

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. As of now Amazon does not provide immutable lock capability for the objects that are transitioned to Glacier using the Amazon lifecycle policy.

GLACIER 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 created using GLACIER 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 storage class (using Amazon S3 services). This results in difference in speed that must be considered when selecting between the two.

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

GLACIER_VAULT storage class

Cost of operation for GLACIER and GLACIER_VAULT storage class is approximately the same with GLACIER being approximately 2% higher than GLACIER_VAULT.

Use GLACIER storage class if you do not plan to use immutable vault lock capability. Amazon Glacier gives you low cost storage, better speed, flexibility, and visibility.

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 182.

Amazon takes 24 hours to refresh archive inventory. Hence, archives uploaded during backup done using GLACIER_VAULT storage class will be 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.

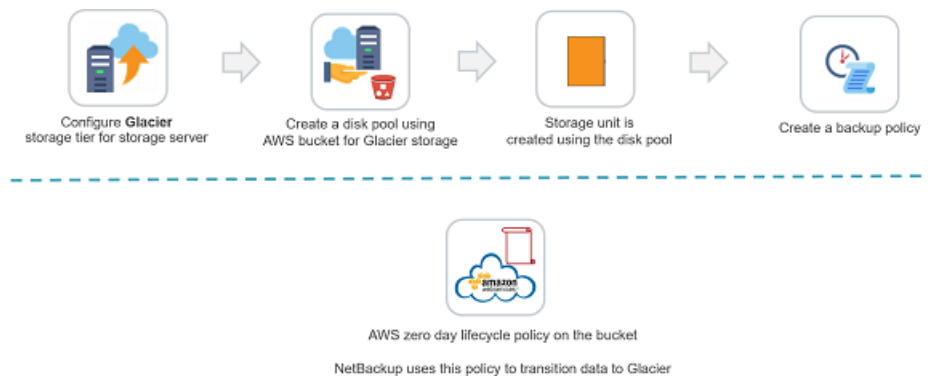
There are architectural differences between GLACIER_VAULT storage class (using Amazon Glacier services) and GLACIER storage class (using Amazon S3 services). This results in difference in speed that must be considered when selecting between the two.

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 storage class. During the backup process, NetBackup internally uses the Amazon's zero-day lifecycle policy to transition data to Glacier. AWS lifecycle policy is a lifecycle rule defined to transition objects to the Glacier storage class in 0 (zero) days after creation. The following diagram illustrates the configuration process:

Figure 2-4 Protecting data in Amazon Glacier



To configure a cloud storage server for Amazon GLACIER storage class

- 1 Configure the Amazon GLACIER cloud storage server.
 See [“Configuring a storage server for cloud storage”](#) on page 100.
- 2 Create a disk pool using the Amazon bucket for GLACIER storage.
 See [“Configuring a disk pool for cloud storage”](#) on page 120.
- 3 Create a backup policy.
 See [“Creating a backup policy”](#) on page 140.
 See the *NetBackup Administrator's Guide, Volume I*:

Best practices

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

- Ensure that Amazon Glacier is supported for the region to which the bucket belongs.

- Ensure that the selected bucket does not have any existing Amazon lifecycle policy.
- 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 data sent to Glacier, there is an inherent time delay of around four hours per fragment of the backup image. For phase 2 of image imports, this time delay is prevalent for image(s) in the Glacier storage. However, if you enable **True Image Recovery** in the policy, time spent for phase 2 imports reduces drastically from four 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, you must backup using multi-streaming that creates multiple images at logical boundaries.
- Workload Granular Recovery (GRT) or VMware Single File Restore (SFR), increases the timeout on the master, 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 Amazon Glacier. Do not select the **Accelerator** check box.
- CloudCatalyst with Glacier is not supported.

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.

About backing up data to Amazon Glacier

When a NetBackup backup job is run to backup data in to Amazon Glacier, NetBackup internally uses the Amazon zero day lifecycle policy. The data objects are tagged as **NetBackupType=LongTerm**. Only the data objects are backed up to Glacier storage, while the metadata objects reside in the Standard storage.

Figure 2-5 Backing up to Amazon Glacier



To duplicate tape data to Amazon Glacier

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

<http://www.veritas.com/docs/DOC5332>

For information on the command, see *Veritas NetBackup Commands Reference Guide*.

About restoring data from Amazon Glacier

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

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

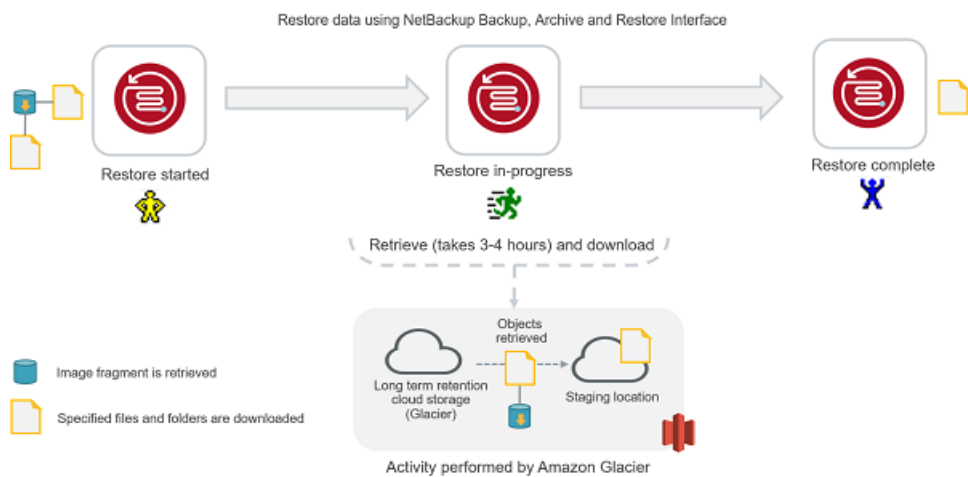
The entire restore operation takes a minimum of 3 – 5 hours. The objects are available at the Amazon staging location depending on the retrieval retention period you have specified. Veritas recommends that you set the retrieval retention period

to a minimum of 3 days. After the retrieval retention period expires, the data is transitioned back to Amazon Glacier.

Note: NetBackup supports Amazon Standard retrievals, which complete within minimum 3 – 5 hours.

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

Figure 2-6 Restoring from Amazon GLACIER




Considerations with Restore of Image Fragments

If the files and folders, you want to restore belong to multiple image fragment 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 will be added to the complete restore time.

Figure 2-7 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 2-8 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 100.

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 120.

Note: If you cannot see the desired vault, it means 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.
 See [“Creating a backup policy”](#) on page 140.

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 two 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 expiry will fail. See “[Handling image cleanup failures for Amazon Glacier vault](#)” on page 181.
- Veritas recommends you use a vault as a secondary target for backing up data.
- If you plan to use the vault lock policy, ensure 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 stored per backup.
- Select True Image Recovery option wherever possible to reduce time and cost for image imports.

To retrieve data sent to Glacier, there is an inherent time delay of around four hours per fragment of the backup image. For phase 2 of image imports, this time delay is prevalent for image(s) in the Glacier storage. However, if you enable **True Image Recovery** in the policy, time spent for phase 2 imports reduces drastically from four 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.
- **CloudCatalyst** with **GLACIER_VAULT** is not supported.
- 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 you have the required S3 related IAM USER permissions.

For issues related to permissions, See [“Troubleshooting failures due to missing Amazon IAM permissions”](#) on page 184.

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 2-9 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.
- Veritas recommends using the LIFECYCLE storage class for moving data from other cloud storage classes to Glacier. See [“Protecting data using Amazon's cloud tiering”](#) on page 45. 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 2-10 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 2-11 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's cloud tiering

Use the LIFECYCLE storage class to protect your data using cloud tiering. Cloud tiering allows you to 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 LIFECYCLE storage class

- 1 Configure the Amazon LIFECYCLE cloud storage server.
 See [“Configuring a storage server for cloud storage”](#) on page 100.
- 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 112.

- 3 Create a disk pool for the LIFECYCLE storage class.
See “[Configuring a disk pool for cloud storage](#)” on page 120.
- 4 Create a backup policy.
See “[Creating a backup policy](#)” on page 140.

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 multi-streaming 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 LIFECYCLE. Do not select the **Accelerator** check box.
- CloudCatalyst with LIFECYCLE is not supported.

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.

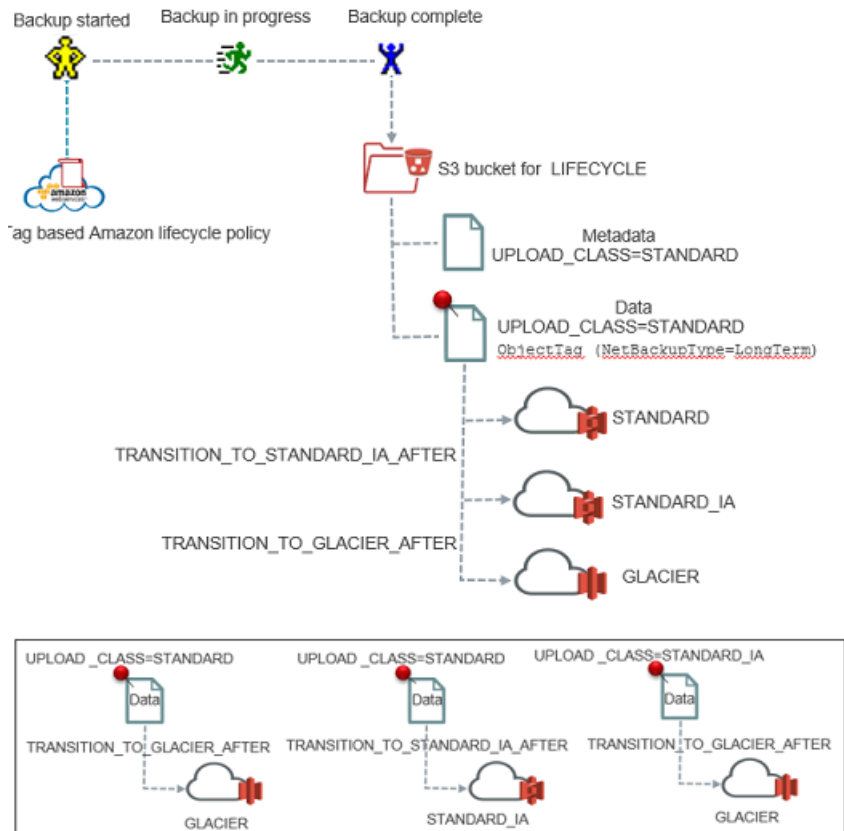
- Also ensure that you also have the required IAM USER permissions.

About backing up data using 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 2-12 Back up process for 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 112.

See [“About protecting data in Amazon Glacier”](#) on page 35.

About restoring data from 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 37.

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.

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 2-10 Microsoft Azure storage API type information and topics

Information	Topic
Certified vendors	See “ Microsoft Azure cloud storage vendors certified for NetBackup ” on page 50.
Requirements	See “ Microsoft Azure storage type requirements ” on page 50.
Storage server configuration options	See “ Microsoft Azure cloud storage provider options ” on page 51.
SSL and proxy options	See “ Microsoft Azure advanced server configuration options ” on page 55.

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

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 8.1.2 release: http://jayne.clab.veritas.com/cl/hcl/nbu/80/out/nbu_80_hcl.html#cloud_storage_solutions

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

See “[About Microsoft Azure cloud storage API type](#)” on page 50.

Microsoft Azure storage type requirements

Table 2-11 describes the details and requirements of Microsoft Azure cloud storage in NetBackup.

Table 2-11 Microsoft Azure cloud storage requirements

Requirement	Details
License requirement	You must have a NetBackup license that allows for cloud storage.

Table 2-11 Microsoft Azure cloud storage requirements (*continued*)

Requirement	Details
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).
Container names	<p>Veritas recommends that you use NetBackup to create the container that you use with NetBackup.</p> <p>The following are the NetBackup requirements for container names:</p> <ul style="list-style-type: none"> ■ 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. <p>You can refer to the following link:</p> <p>https://msdn.microsoft.com/en-us/library/azure/dd135715.aspx</p>

See “[About Microsoft Azure cloud storage API type](#)” on page 50.

Microsoft Azure cloud storage provider options

[Figure 2-13](#) shows the **Cloud Storage Configuration Wizard** panel for Microsoft Azure cloud storage.

Figure 2-13 Cloud Storage Server Configuration Wizard panel for Microsoft Azure

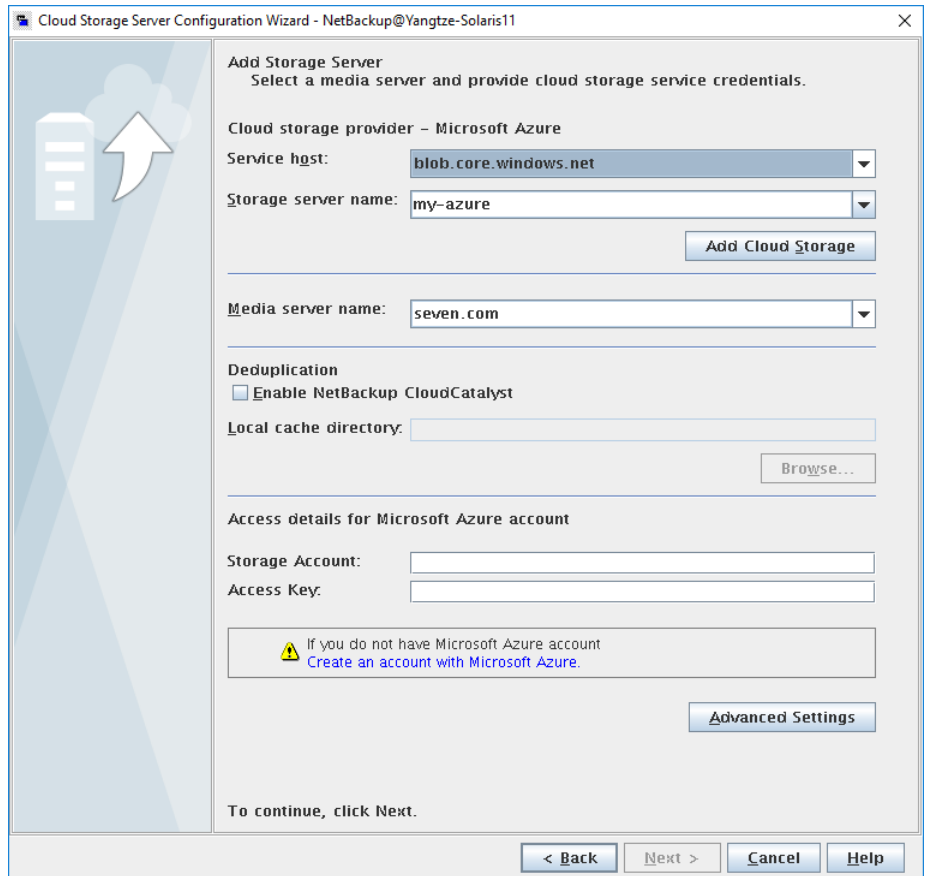


Table 2-12 describes the storage server configuration options for Microsoft Azure.

Table 2-12 Microsoft Azure storage server configuration options

Field name	Required content
Service host	<p>Service host is the host name of the cloud service end point of Microsoft Azure.</p> <p>The Service host drop-down list displays part of the service host URL that also comprises Storage Account.</p> <p>Example of a service host URL:</p> <p><i>storage_account.blob.core.windows.net</i></p> <p>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.</p>
Storage server name	<p>Displays the default Azure storage server, which is my-azure. You can select a storage server other than the default one.</p> <p>The drop-down list displays only those names that are available for use.</p> <p>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.</p> <p>Note: Veritas recommends 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.</p>

Table 2-12 Microsoft Azure storage server configuration options (*continued*)

Field name	Required content
Deduplication	<p>Enabling this option creates a CloudCatalyst storage server that can be used to upload deduplicated data to the cloud.</p> <p>This option is grayed out if any of the following cases are true:</p> <ul style="list-style-type: none"> ■ The selected media server does not have NetBackup 8.1 or later installed. ■ CloudCatalyst does not support the media server operating system. ■ CloudCatalyst does not support the cloud vendor. <p>See the NetBackup compatibility lists for support information: http://www.netbackup.com/compatibility</p> <p>For information about CloudCatalyst, see the <i>NetBackup Deduplication Guide</i>: http://www.veritas.com/docs/DOC5332</p>
Local cache directory	<p>Enter the mount path to be used as the storage path on the CloudCatalyst storage server.</p> <p>For example: <code>/space/mnt/esfs</code></p> <p>The deduplicated data is written to this local cache directory before it is uploaded to the cloud. The larger the cache, the more likely that NetBackup can service requests locally, avoiding cloud access to read and write.</p> <p>Notes:</p> <ul style="list-style-type: none"> ■ This path should be to a file system which is dedicated for CloudCatalyst cache use. Inaccurate cache eviction occurs if the path shares any storage with other data or applications. ■ NetBackup manages the files in the local cache directory. Users should not manually delete files in this directory.
Media server name	<p>Select a NetBackup media server from the drop-down list.</p> <p>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:</p> <p>See “About the NetBackup media servers for cloud storage” on page 97.</p> <p>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.</p>

Table 2-12 Microsoft Azure storage server configuration options (*continued*)

Field name	Required content
Storage Account	<p>Enter the storage account that you want to use for your cloud backups.</p> <p>For more information about Microsoft Azure storage service, refer to the Microsoft Azure documentation.</p> <p>http://azure.microsoft.com</p> <p>Create the storage account using the following URL:</p> <p>https://portal.azure.com</p>
Access key	<p>Enter your Azure access key. You can enter the primary access key or the secondary access key. It must be 100 or fewer characters.</p> <p>Refer to the following URL for the access key:</p> <p>https://portal.azure.com</p>
Advanced Settings	<p>To change SSL or proxy settings for Azure, click Advanced Settings.</p> <p>See “Microsoft Azure advanced server configuration options” on page 55.</p>
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 ARCHIEVE	Select ARCHIVE option for long term retention.

See “[About Microsoft Azure cloud storage API type](#)” on page 50.

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 2-13 General settings options

Option	Description
Use SSL	<p>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.</p> <ul style="list-style-type: none"> ■ 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. <p>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.</p>

Table 2-14 Proxy Settings tab options

Option	Description
Use Proxy Server	<p>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:</p> <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ HTTP Note: You need to provide the proxy credentials for HTTP proxy type. ■ SOCKS ■ SOCKS4 ■ SOCKS5 ■ SOCKS4A
Use Proxy Tunneling	<p>You can enable proxy tunneling for HTTP proxy type.</p> <p>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.</p> <p>The data passes through the proxy server without reading the headers or data from the connection.</p>

Table 2-14 Proxy Settings tab options (*continued*)

Option	Description
Authentication Type	<p>You can select one of the following authentication types if you are using HTTP proxy type.</p> <ul style="list-style-type: none"> ■ None— Authentication is not enabled. Username and password is not required. ■ NTLM—Username and password needed. ■ Basic—Username and password needed. <p>Username is the username of the proxy server</p> <p>Password can be empty. You can use maximum 256 characters.</p>

See [“About Microsoft Azure cloud storage API type”](#) on page 50.

About OpenStack Swift cloud storage API type

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

Table 2-15 OpenStack Swift storage API type information and topics

Information	Topic
Certified vendors	See “OpenStack Swift cloud storage vendors certified for NetBackup” on page 58.
Requirements	See “OpenStack Swift storage type requirements” on page 58.
Storage server configuration options	See “OpenStack Swift cloud storage provider options” on page 59.
Region and host configuration options	See “OpenStack Swift storage region options” on page 63.
Cloud instance configuration options	See “OpenStack Swift add cloud storage configuration options” on page 65.
Proxy connection options	See “OpenStack Swift proxy settings” on page 65.

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

OpenStack Swift cloud storage vendors certified for NetBackup

Click the following link to identify the vendors who are certified for NetBackup cloud storage using the OpenStack Swift storage API as of the NetBackup 8.1.2 release:
http://jayne.clab.veritas.com/cl/hcl/nbu/80/out/nbu_80_hcl.html#cloud_storage_solutions

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

See “[About OpenStack Swift cloud storage API type](#)” on page 57.

OpenStack Swift storage type requirements

The following table provides links to the details and requirements of OpenStack Swift compatible cloud.

Table 2-16 OpenStack Swift compatible cloud storage requirements

Requirement	Details
License requirement	You must have a NetBackup license that allows for cloud storage.
Storage account requirements	<p>You must obtain the credentials required to access the cloud storage account.</p> <p>If you use authentication V1, only the user name and password are required to validate the user to access the cloud storage.</p> <p>If you use authentication version Identity V2, the user name, password, and either tenant ID or tenant name is required to validate the user to access the cloud storage.</p>

Table 2-16 OpenStack Swift compatible cloud storage requirements
(continued)

Requirement	Details
Containers	<p>The containers for OpenStack Swift compliant cloud providers cannot be created in NetBackup. You must use the native cloud tools to create a container.</p> <p>The container names must conform to the following requirements:</p> <ul style="list-style-type: none"> ■ The container 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. ■ Any of the following characters (you cannot use these as the first character in the container name): Period (.), underscore (_), and dash (-). <i>Exception:</i> If you use SSL for communication, you cannot use a period. By default, NetBackup uses SSL for communication. See “NetBackup cloud storage server connection properties” on page 112. <p>Note: Only those containers are listed in NetBackup that follow these naming conventions.</p>

See [“About OpenStack Swift cloud storage API type”](#) on page 57.

OpenStack Swift cloud storage provider options

[Figure 2-14](#) shows the cloud storage provider wizard panel for OpenStack Swift-compliant cloud storage. The panel includes cloud provider and access information.

Figure 2-14 Cloud Storage Server Configuration Wizard panel




Table 2-17 describes configuration options for OpenStack Swift cloud storage.

Table 2-17 OpenStack Swift provider and access details

Field name	Required content
Cloud storage provider	Displays the name of the selected cloud provider.
Cloud storage name	Select the cloud storage name from the list. If the list is empty, you must add a cloud storage instance. See the Add Cloud Storage option description.
Add Cloud Storage	Click the add cloud storage option, then add, select, or enter the required information. See “OpenStack Swift add cloud storage configuration options” on page 65.

Table 2-17 OpenStack Swift provider and access details (*continued*)

Field name	Required content
Tenant ID / Tenant Name	<p>Based on the selection, enter either the tenant ID or tenant name that is associated with your cloud storage credentials.</p> <p>Note: This field is visible only if you selected the Identity v2 Authentication version in the Add Cloud Storage dialog box.</p> <p>See “OpenStack Swift add cloud storage configuration options” on page 65.</p>
User name	<p>Enter the user name that is required to access the cloud storage.</p>
Password	<p>Enter the password that is required to access the cloud storage. It must be 100 or fewer characters.</p>
Proxy Settings	<p>To change the default storage server for your cloud vendor or specify the maximum number of network connections, click Advanced Settings.</p>
User ID	<p>Based on the selection, enter either the User ID or the User Name that is associated with your cloud storage credentials. When you provide User ID, User Name and Domain information is not required.</p> <p>Note: This field is visible only if you selected the Identity v3 Authentication version in the Authentication version dialog box.</p> <p>See “OpenStack Swift add cloud storage configuration options” on page 65.</p>
Domain ID / Domain name (for user details)	<p>Based on the selection, enter either the user’s Domain ID or Domain Name that is associated with your cloud storage credentials.</p> <p>Note: This field is visible only if you selected the Identity v3 Authentication version in the Authentication version dialog box.</p> <p>See “OpenStack Swift add cloud storage configuration options” on page 65.</p>
Project ID / Project Name	<p>Based on the selection, enter either the Project ID or Project Name that is associated with your cloud storage credentials. When you provide Project ID, Project Name and Domain information is not required.</p> <p>Note: This field is visible only if you selected the Identity v3 Authentication version in the Authentication version dialog box.</p> <p>See “OpenStack Swift add cloud storage configuration options” on page 65.</p>

Table 2-17 OpenStack Swift provider and access details (*continued*)

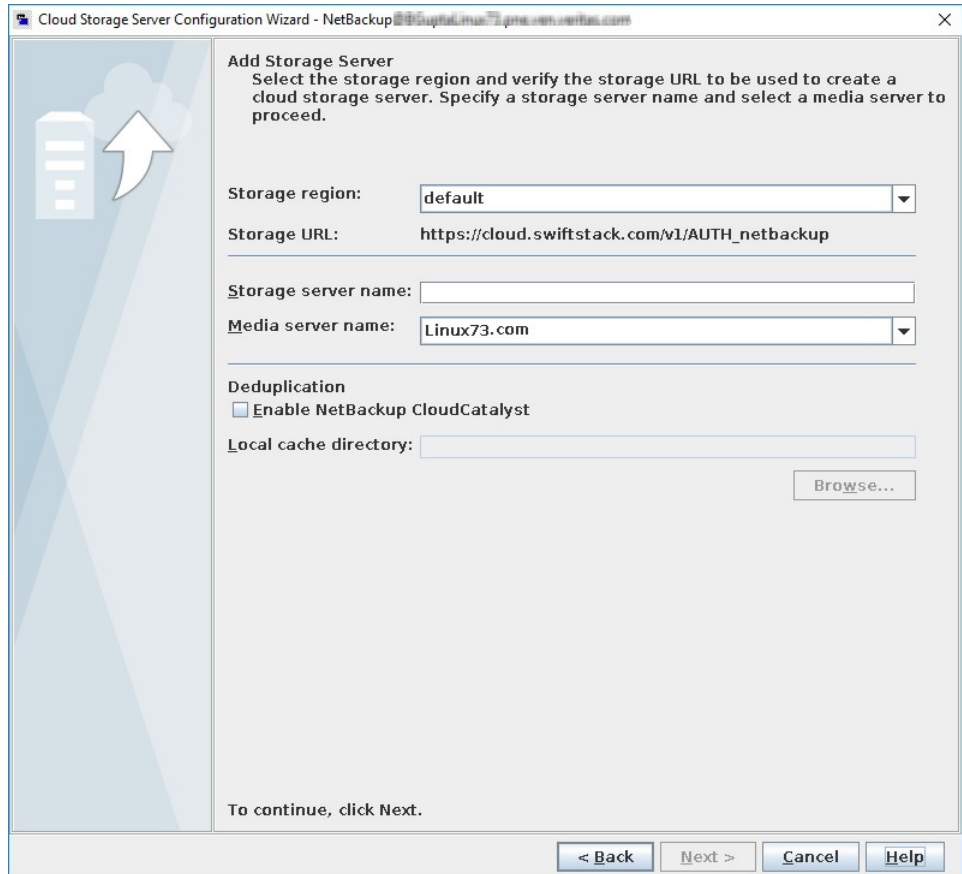
Field name	Required content
Domain ID / Domain name (for project details)	<p>Based on the selection, enter either the project's Domain ID or Domain Name that is associated with your cloud storage credentials.</p> <p>Note: This field is visible only if you selected the Identity v3 Authentication version in the Authentication version dialog box.</p> <p>See “OpenStack Swift add cloud storage configuration options” on page 65.</p>
Deduplication	<p>Enabling this option creates a CloudCatalyst storage server that can be used to upload deduplicated data to the cloud.</p> <p>This option is grayed out if any of the following cases are true:</p> <ul style="list-style-type: none"> ■ The selected media server does not have NetBackup 8.1 or later installed. ■ CloudCatalyst does not support the media server operating system. ■ CloudCatalyst does not support the cloud vendor. <p>See the NetBackup compatibility lists for support information: http://www.netbackup.com/compatibility</p> <p>For information about CloudCatalyst, see the <i>NetBackup Deduplication Guide</i>: http://www.veritas.com/docs/DOC5332</p>
Local cache directory	<p>Enter the mount path to be used as the storage path on the CloudCatalyst storage server.</p> <p>For example: <code>/space/mnt/esfs</code></p> <p>The deduplicated data is written to this local cache directory before it is uploaded to the cloud. The larger the cache, the more likely that NetBackup can service requests locally, avoiding cloud access to read and write.</p> <p>Notes:</p> <ul style="list-style-type: none"> ■ This path should be to a file system which is dedicated for CloudCatalyst cache use. Inaccurate cache eviction occurs if the path shares any storage with other data or applications. ■ NetBackup manages the files in the local cache directory. Users should not manually delete files in this directory.

See [“About OpenStack Swift cloud storage API type”](#) on page 57.

OpenStack Swift storage region options

Figure 2-15 shows the storage region wizard panel for OpenStack Swift-compliant cloud storage. The panel includes storage region and storage host information.

Figure 2-15 Cloud Storage Server Configuration Wizard panel



Provider and access details are used to map the cloud storage settings to NetBackup storage settings. The cloud storage region is mapped to the NetBackup storage server. All the backups that are targeted to the NetBackup storage server use the cloud storage region to which it is mapped.

Note: One cloud storage region is mapped to one NetBackup storage server.

Table 2-18 describes configuration options for OpenStack Swift cloud storage.

Table 2-18 OpenStack Swift region and host details

Field name	Description
Storage region	<p>Select the cloud storage region.</p> <p>You may use the cloud storage region that is geographically closest to the NetBackup media server that sends the backups to the cloud. Contact your storage administrator for more details.</p> <p>Note: This field is visible only if you selected the Identity v2 Authentication version in the Add Cloud Storage dialog box.</p> <p>See “OpenStack Swift add cloud storage configuration options” on page 65.</p>
Storage URL	<p>The cloud storage URL is auto-populated based on the storage region selection. This field is non-editable and is only for your reference.</p> <p>Note: This field is visible only if you selected the Identity v2 Authentication version in the Add Cloud Storage dialog box.</p> <p>See “OpenStack Swift add cloud storage configuration options” on page 65.</p>
Storage server name	<p>Enter a unique name for the storage server.</p> <p>Note: Veritas recommends that a storage server name that you add while configuring an OpenStack Swift compatible cloud provider should be a logical name and should not match a physical host name. For example: When you add an Oracle storage server, avoid using names like ‘oracle.com’ or ‘oracle123.com’. These servers may be physical hosts, which can cause failures during cloud storage configuration. Instead, use storage server names like ‘oracle1’ or ‘oracleserver1’ and so on.</p>
Media server name	<p>Select a NetBackup media server from the drop-down list. The drop-down list displays only NetBackup 8.1.2 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:</p> <p>See “About the NetBackup media servers for cloud storage” on page 97.</p> <p>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.</p>

See [“About OpenStack Swift cloud storage API type”](#) on page 57.

OpenStack Swift add cloud storage configuration options

The following table describes the configuration options for the **Add Cloud Storage** dialog box. It appears when you click **Add Cloud Storage** on the wizard panel for OpenStack providers.

Table 2-19 Add Cloud Storage

Field	Description
Cloud storage provider	The cloud storage provider from the previous wizard panel is displayed.
Cloud storage name	Enter a unique name to identify the authentication service endpoint. You can reuse the same authentication service endpoint for another storage server.
Authentication location	This field is not visible for cloud providers with custom authentication URLs. Select the authentication location of the cloud storage, otherwise, select Other . Note: If you select Other , you must enter the authentication URL.
Authentication version	Select the authentication version that you want to use. Select Do not use identity service if you do not want to authenticate using the OpenStack's Identity APIs.
Authentication URL	Enter the authentication URL that your cloud vendor provided. Authentication URL comprises of either HTTP or HTTPS and port number. For example, <code>http://mycloud.example.com:5000/v2.0/tokens</code> For custom instance, to use IPv6 endpoint, you must update or create a new instance with the IPv6 equivalent authentication URL.

See [“OpenStack Swift cloud storage provider options”](#) on page 59.

OpenStack Swift proxy settings

For security purpose, you can use a proxy server to establish communication with the cloud storage.

The following table describes the options of the **Proxy Settings** dialog box.

Table 2-20 Proxy settings for OpenStack Swift

Option	Description
Use Proxy Server	<p>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:</p> <ul style="list-style-type: none"> ■ Proxy Host—Specify IP address or name of the proxy server. ■ Proxy Port—Specify port number of the proxy server. Possible values: 1-65535 ■ Proxy Type— You can select one of the following proxy types: <ul style="list-style-type: none"> ■ HTTP <p>Note: You need to provide the proxy credentials for HTTP proxy type.</p> <ul style="list-style-type: none"> ■ SOCKS ■ SOCKS4 ■ SOCKS5 ■ SOCKS4A
Use Proxy Tunneling	<p>You can enable proxy tunneling for HTTP proxy type.</p> <p>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.</p> <p>The data passes through the proxy server without reading the headers or data from the connection.</p>
Authentication Type	<p>You can select one of the following authentication types if you are using HTTP proxy type.</p> <ul style="list-style-type: none"> ■ None— Authentication is not enabled. Username and password is not required. ■ NTLM—Username and password needed. ■ Basic—Username and password needed. <p>Username is the username of the proxy server</p> <p>Password can be empty. You can use maximum 256 characters.</p>

See [“About OpenStack Swift cloud storage API type”](#) on page 57.

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 key management 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](#)
- [Changing cloud storage server properties](#)

- [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](#)
- [Enabling NetBackup Accelerator with cloud storage](#)
- [Enabling optimized synthetic backups with cloud storage](#)
- [Creating a backup policy](#)
- [Changing cloud storage disk pool properties](#)
- [Managing Certification Authorities \(CA\) for NetBackup Cloud](#)

Before you begin to configure cloud storage in NetBackup

Veritas recommends 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 Veritas 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: Veritas 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:

<http://www.veritas.com/docs/DOC5332>

<http://www.veritas.com/docs/DOC5332>

Table 3-1 Overview of the NetBackup cloud configuration process

Step	Task	More information
Step 1	Create NetBackup log file directories on the master server and the media servers	See "NetBackup cloud storage log files" on page 163. See "Creating NetBackup log file directories for cloud storage" on page 163.
Step 2	Review the cloud installation requirements	See "Cloud installation requirements" on page 71.
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 71.
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 76.
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.

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

Step	Task	More information
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 88.
Step 8	Understand key management for encryption	Encryption is optional. See “ About data encryption for cloud storage ” on page 92. See “ About key management 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 78. See “ Configuring a storage server for cloud storage ” on page 100.
Step 10	Configure the disk pool	See “ About cloud storage disk pools ” on page 120. See “ Configuring a disk pool for cloud storage ” on page 120.
Step 11	Configure additional storage server properties	See “ NetBackup cloud storage server properties ” on page 108. See “ Changing cloud storage server properties ” on page 106.
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 131.
Step 13	Configure a storage unit	See “ Configuring a storage unit for cloud storage ” on page 132.
Step 14	Configure NetBackup Accelerator and optimized synthetic backups	Accelerator and optimized synthetic backups are optional. See “ About NetBackup Accelerator and NetBackup Optimized Synthetic backups ” on page 136. See “ Enabling NetBackup Accelerator with cloud storage ” on page 137. See “ Changing cloud storage server properties ” on page 106.
Step 15	Configure a backup policy	See “ Creating a backup policy ” on page 140. See the <i>NetBackup Administrator's Guide, Volume I</i> : http://www.veritas.com/docs/DOC5332

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	<p>For the operating systems that NetBackup supports for cloud storage, see the NetBackup operating system compatibility list available through the following URL:</p> <p>http://www.netbackup.com/compatibility</p> <p>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.</p>
Cloud storage provider account	<p>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.</p> <p>You can create this account in the Cloud Storage Configuration Wizard.</p> <p>See “About the cloud storage vendors for NetBackup” on page 14.</p>
NetBackup cloud storage licensing	<p>NetBackup cloud storage is licensed separately from base NetBackup.</p> <p>The license also enables the Use Accelerator feature on the NetBackup policy Attributes tab. Accelerator increases the speed of full backups for files systems.</p>

Scalable Storage properties

The **Scalable Storage Cloud Settings** properties contain information about encryption, metering, bandwidth throttling, and network connections between the NetBackup hosts and your cloud storage provider.

The **Scalable Storage** properties appear only if the host is supported for cloud storage. See the NetBackup hardware compatibility list for your release available through the following URL:

<http://www.netbackup.com/compatibility>

<http://www.netbackup.com/compatibility>

The **Scalable Storage** properties apply to currently selected media servers.

Figure 3-1 Scalable Storage Cloud Settings host properties




Table 3-3 describes the properties.

Table 3-3 Scalable Storage Cloud Settings host properties

Property	Description
Key Management Server (KMS) Name	If you configured the NetBackup Key Management Service (KMS), the name of the KMS server.
Metering Interval	Determines how often NetBackup gathers connection information for reporting purposes. NetBackup OpsCenter uses the information that is collected to create reports. The value is set in seconds. The default setting is 300 seconds (5 minutes). If this value 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.

Table 3-3 Scalable Storage Cloud Settings host properties (*continued*)

Property	Description
Advanced Settings	Click Advanced Settings to specify additional settings for throttling. See “ Configuring advanced bandwidth throttling settings ” on page 73. See “ Advanced bandwidth throttling settings ” on page 74.
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. If the media server is not a CloudCatalyst storage server, a value over 100 is generally not needed. If the media server is a CloudCatalyst storage server, change the value to 160 or more.

<http://www.veritas.com/docs/DOC5332>

<http://www.veritas.com/docs/DOC5332>

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.

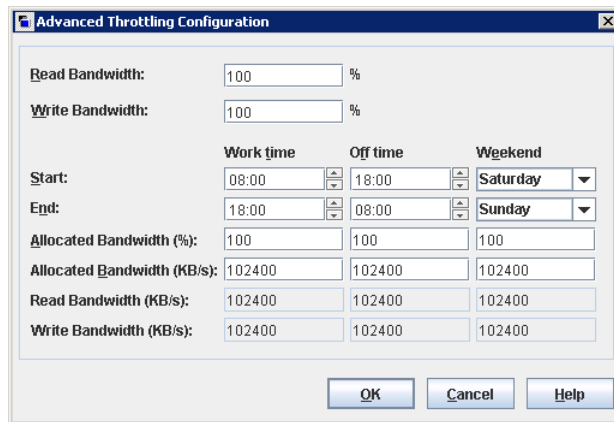
The total bandwidth and the bandwidth sampling interval are configured on the **Cloud Settings** tab of the **Scalable Storage** host properties screen.

See “[Scalable Storage properties](#)” on page 71.

To configure advanced bandwidth throttling settings

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Media Servers** in the left pane.
- 2 In the right pane, select the host on which to specify properties.
- 3 Click **Actions > Properties**.
- 4 In the properties dialog box left pane, select **Scalable Storage**.
- 5 In the right pane, click **Advanced Settings**. The **Advanced Throttling Configuration** dialog box appears.

The following is an example of the dialog box:



- 6 Configure the settings and then click **OK**.
See [“Advanced bandwidth throttling settings”](#) on page 74.

Advanced bandwidth throttling settings

The following table describes the advanced bandwidth throttling settings.

Table 3-4 Advanced Throttling Configuration settings

Property	Description
Read Bandwidth	<p>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.</p> <p>If there is insufficient bandwidth to transmit the specified amount of data within a few minutes, restore or replication failures may occur due to timeouts.</p> <p>Consider the total load of simultaneous jobs on multiple media servers when you calculate the required bandwidth.</p> <p>Default value: 100</p> <p>Possible values: 0 to 100</p>
Write Bandwidth	<p>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.</p> <p>If there is insufficient bandwidth to transmit the specified amount of data within a few minutes, backup failures may occur due to timeouts.</p> <p>Consider the total load of simultaneous jobs on multiple media servers when you calculate the required bandwidth.</p> <p>Default value: 100</p> <p>Possible values: 0 to 100</p>
Work time	<p>Use this field to specify the time interval that is considered work time for the cloud connection.</p> <p>Specify a start time and end time in 24-hour format. For example, 2:00 P.M. is 14:00.</p> <p>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.</p>

Table 3-4 Advanced Throttling Configuration settings (*continued*)

Property	Description
Off time	<p>Use this field to specify the time interval that is considered off time for the cloud connection.</p> <p>Specify a start time and end time in 24-hour format. For example, 2:00 P.M. is 14:00.</p> <p>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.</p>
Weekend	<p>Specify the start and stop time for the weekend.</p> <p>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.</p>
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.
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.

See [“Scalable Storage properties”](#) on page 71.

See [“Configuring advanced bandwidth throttling settings”](#) on page 73.

Cloud Storage properties

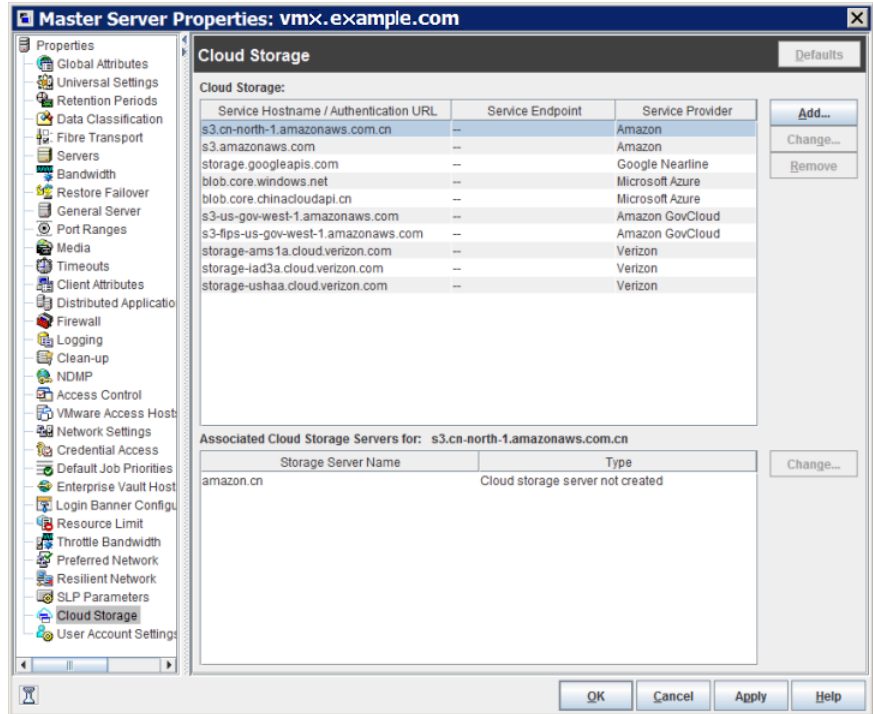
The NetBackup **Cloud Storage** properties in the **NetBackup Administration Console** apply to the currently selected master server.

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 Hostname** is available or required.

Figure 3-2 Cloud Storage host properties



Cloud Storage host properties contain the following properties:

Table 3-5 Cloud Storage

Property	Description
<p>Cloud Storage</p>	<p>The cloud storage that corresponds to the various cloud service providers that NetBackup supports are listed here.</p> <p>See “Amazon S3 cloud storage options” on page 23.</p> <p>See “Adding a cloud storage instance” on page 78.</p> <p>See “Amazon S3 cloud storage options” on page 23.</p> <p>See “Changing cloud storage host properties” on page 79.</p> <p>See “Deleting a cloud storage host instance” on page 80.</p>

Table 3-5 Cloud Storage (*continued*)

Property	Description
Associated Storage Servers for	The cloud storage servers that correspond to the selected cloud storage are displayed. See “Amazon S3 advanced server configuration options” on page 25. See “Amazon S3 credentials broker details” on page 28. See “Changing cloud storage host properties” on page 79.

Note: Changes that you make in the **Cloud Storage** dialog box are applied before you click **OK** in the **Host Properties** dialog box.

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

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. A custom cloud storage instance appears in the **Cloud Storage Server Configuration Wizard** when you configure a storage server.

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:

By using NetBackup
Master Server Properties

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 Cloud Storage host properties”](#) on page 79.

By using the **Cloud Storage Server Configuration Wizard**

With this method, you add the instance at the same time as when you configure the storage server in NetBackup. The wizard that configures the storage is *not* populated with the instance details until you add them in the wizard itself.

See [“Configuring a storage server for cloud storage”](#) on page 100.

See [“Cloud Storage properties”](#) on page 76.

To add a cloud storage instance in Cloud Storage host properties

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management** > **Host Properties** > **Master Servers** in the left pane.
- 2 In the right pane, select the master server on which to add the cloud storage instance.
- 3 On the **Actions** menu, click **Properties**.
- 4 In the properties dialog box left pane, select **Cloud Storage**.
- 5 In the right pane, click **Add**.
- 6 In the **Add Cloud Storage** dialog box, configure the settings.
See [“Amazon S3 cloud storage options”](#) on page 23.
- 7 After you configure the settings, click **OK**.

Changing cloud storage host properties

From the **Cloud Storage Master Server Properties**, you can change the following properties:

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 cloud storage host properties” on page 79.
Associated cloud storage server properties	See “To change associated cloud storage server host properties” on page 80.

See [“Cloud Storage properties”](#) on page 76.

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

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

To change cloud storage host properties

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management** > **Host Properties** > **Master Servers** in the left pane.
- 2 In the right pane, select the master server on which to specify properties.
- 3 On the **Actions** menu, click **Properties**.
- 4 In the left pane of the **Master Server Properties** dialog box, select **Cloud Storage**.

5 In the **Cloud Storage** list in the right pane, select the wanted cloud storage.

6 Click **Change** adjacent to the **Cloud Storage** list.

7 In the **Change Cloud Storage** dialog box, change the properties.

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

8 Click **OK** in the **Change Cloud Storage** dialog box.

9 Click **OK** to close the **Master Server Properties** dialog box.

To change associated cloud storage server host properties

1 In the **NetBackup Administration Console**, expand **NetBackup Management** > **Host Properties** > **Master Servers** in the left pane.

2 In the right pane, select the master server on which to specify properties.

3 On the **Actions** menu, click **Properties**.

4 In the left pane of the **Master Server Properties** dialog box, select **Cloud Storage**.

5 In the **Associated Cloud Storage Servers for** list in the right pane, select the wanted storage server.

6 Click **Change** adjacent to the **Associated Cloud Storage Servers for** list.

7 In the **Cloud Storage Server Configuration** dialog box, change the properties.

See [“Amazon S3 advanced server configuration options”](#) on page 25.

See [“Amazon S3 credentials broker details”](#) on page 28.

8 Click **OK** in the **Change Cloud Storage** dialog box.

9 Click **OK** to close the **Master Server Properties** dialog box.

Deleting a cloud storage host instance

You can delete your custom cloud storage (cloud instance) by using the **Cloud Storage Master Server Properties**. You cannot delete the cloud storage instances that were delivered with NetBackup.

See [“Cloud Storage properties”](#) on page 76.

To delete a cloud storage host instance

1 In the **NetBackup Administration Console**, expand **NetBackup Management** > **Host Properties** > **Master Servers** in the left pane.

2 In the right pane, select the master server on which to specify properties.

3 On the **Actions** menu, click **Properties**.

- 4 In the left pane of the **Master Server Properties** dialog box, select **Cloud Storage**.
- 5 In the **Cloud Storage** list in the right pane, select the wanted cloud storage.
- 6 Click **Remove**.
- 7 In the **Remove the Cloud Storage** dialog box, click **Yes**.
- 8 Click **OK** to close the **Master Server Properties** dialog box.

About the NetBackup CloudStore Service Container

This 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 older media servers that are configured for cloud storage.

This container hosts the throttling service and the metering data collector service. NetBackup OpsCenter uses the metering data for monitoring and reporting.

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

See [“Scalable Storage properties”](#) on page 71.

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

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 84.

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

Note: For NetBackup releases beyond 8.1.2, 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 *Veritas 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 90.

If the NetBackup master 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:

<http://www.veritas.com/docs/DOC5332>

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 88.

See “[Deploying host ID-based certificates](#)” on page 90.

If the NetBackup master 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 NetBackup 8.0 or later:

<http://www.veritas.com/docs/DOC5332>

NetBackup 7.7 and 7.7.x

The NetBackup hosts that run the CloudStore Service Container require a host name-based certificate. You must use a command to install it on a media server.

See “[Deploying host name-based certificates](#)” on page 88.

Note: You do not need to generate a security certificate, if you have already generated it before configuring the cloud storage.

The host name-based security certificates expire after one year. NetBackup automatically replaces existing certificates with new ones as needed.

Note: The security certificates that are provisioned for other NetBackup features or purposes satisfy the certificate requirement for the NetBackup CloudStore Service Container. The NetBackup Access Control feature uses security certificates, and the NetBackup Administration Console requires security certificates for interhost communication.

If the NetBackup master server is clustered, you must ensure that the active node and the passive node have host named-based certificates. See the 7.7.x version of the *NetBackup Security and Encryption Guide*

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/openv/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 master 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 master 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/wmc/cloud`
On media server versions 7.7.x to 8.1.2, the path is `/usr/opensv/netbackup/db/cloud`.
- **Windows:** `install_path\Veritas\NetBackup\var\global\wmc\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	Veritas recommends that you do not modify this value. Specifies the version of <code>cloudstore.conf</code> file. The default value is 2.
CSSC_PLUGIN_PATH	Veritas recommends that you do not modify this value. Specifies the path where NetBackup cloud storage plug-ins are installed. The default path is as follows: On Windows: <i>install_path</i> \Veritas\NetBackup\bin\ost-plugins On UNIX: <code>/usr/opensv/lib/ost-plugins</code>
CSSC_PORT	This 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 master server fails if the older media servers use a different port.
CSSC_LOG_DIR	Specifies the directory path where <code>csconfig</code> , <code>nbclutil</code> , and cloud plugins generate log files. The default path is as follows: On Windows: <i>install_path</i> \Veritas\NetBackup\logs\nbcssc 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 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.

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

Parameter	Description
<code>CSCONFIG_LOG_FILE</code>	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.
<code>CSSC_IS_SECURE</code>	Specifies if the <code>nbcssc</code> service runs in secure (value 64) or non-secure mode (value 0). The default value is 64.
<code>CSSC_CIPHER_LIST</code>	<p>Specifies the cipher list that NetBackup uses for the following purpose:</p> <ul style="list-style-type: none"> ■ The cloud master host's cipher is used for communication with the cloud service provider. ■ The media server cipher is used for communicating with the cloud master host's <code>nbwmc</code> service and with the cloud service provider. <p>Veritas recommends that you do not modify this value. However, if you want to customize the cipher list, depending on the purpose, you must modify the cipher list in the <code>cloudstore.conf</code> on the master server and the media servers.</p> <p>Note: If the cipher list is invalid, the customized cipher list is replaced by the default cipher list.</p> <p>The default value is <code>AES : !aNULL : @STRENGTH</code>.</p>
<code>CSSC_LOG_LEVEL</code>	<p>Specifies the log level for <code>csconfig</code> and <code>nbclidutil</code> CLI utility logging. Value 0 indicates that the logging is disabled and non-zero value indicates that the logging is enabled.</p> <p>The default value is 0.</p>
<code>CSSC_MASTER_PORT</code>	<p>This is applicable for media server versions 7.7.x to 8.1.2 only. It is not applicable for NetBackup master and media server versions 8.2 and later.</p> <p>This parameter value must be set to 5637.</p> <p>This port is used to provide back-level media server support for older media servers that are configured for cloud storage. Ensure that the older media servers use this port. Communication with the master server fails if the older media servers use a different port.</p>

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

Parameter	Description
CSSC_MASTER_NAME	<p>Specifies the NetBackup master 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 that reside at the following location:</p> <ul style="list-style-type: none"> ■ On Windows: <ul style="list-style-type: none"> <code>install_path\NetBackup\var\global\wmc\cloud</code> On media server versions 7.7.x to 8.1.2, the path is <code>install_path\NetBackup\db\cloud</code>. ■ On UNIX: <ul style="list-style-type: none"> <code>/usr/opensv/var/global/wmc/cloud</code> On media server versions 7.7.x to 8.1.2, the path is <code>/usr/opensv/netbackup/db/cloud</code>.
CSSC_LEGACY_AUTH_ENABLED	<p>Specifies if the <code>nbcssc</code> service has the legacy authentication enabled (value 1) or disabled (0). The default value is 0.</p> <p>Note: Starting from NetBackup 8.1, the <code>CSSC_LEGACY_AUTH_ENABLED</code> option is deprecated. To communicate with legacy media servers, use the Enable insecure communication with 8.0 and earlier hosts option on the NetBackup master server. The option is available in the NetBackup Administration Console on the Security Management > Global Security Settings > Secure Communication tab.</p>

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

Parameter	Description
CSSC_ALLOW_LEGACY_AUTH	<p>Specifies if the master server can communicate with legacy media servers that are configured for cloud storage. Only media server versions 7.7.x to 8.1.2 are supported.</p> <p>The value 1 (default value) indicates that the communication is enabled while the value 0 means that the communication is disabled.</p> <p>Use this parameter in conjunction with the Enable insecure communication with 8.0 and earlier hosts option available in the NetBackup Administration Console GUI (Security Management > Global Security Settings > Security Communication tab).</p> <p>The GUI option allows you to enable or disable master 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 master server communication with back-level cloud storage media servers explicitly.</p> <p>For example, if the GUI option is enabled (default value) and this parameter value is set to 0, the NetBackup master server continues to work with supported back-level media servers as other storage servers. However, legacy cloud storage media servers that use the older method of communication using hardcoded credentials are blocked altogether, thus increasing the security of your NetBackup environment.</p> <p>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 (<code>nbwmc</code>) service for the changes to take effect.</p>

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 master server, depending on your environment. Specify the name of an individual media server or specify `-AllMediaServers`.

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

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

NetBackup appliance (as a NetBackupCLI user): `bpbaz -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 master server. To do so, run the following NetBackup `bpclient` command on the master 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-master host may require an authorization token before it can obtain a host ID-based certificate from the Certificate Authority (master 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-master 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-master host to establish that the master server can be trusted:

```
nbcertcmd -getCACertificate
```

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

```
nbcertcmd -getCertificate
```

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

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

```
nbcertcmd -getCertificate -server master_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-master host to establish that the master server can be trusted:

```
nbcertcmd -getCACertificate
```

- 3 Run the following command on the non-master 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 master 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 the **Cloud Storage Server Configuration Wizard**.

See [“Configuring a storage server for cloud storage”](#) on page 100.

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 media servers that are older than the 7.7.3 version do not support data compression. Therefore, if you have selected an older media server while you configure the cloud storage server, the compression option does not appear on the **Cloud Storage Server Configuration Wizard**.
- 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 150.
- 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, Veritas recommends not to enable compression for backing up incompressible data such as policy data and so on.
- Veritas recommends not 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. The NetBackup **Cloud Storage Server Configuration Wizard** and the **Disk Pool Configuration Wizard** include the steps that configure key management and encryption.

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

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

More information about data-at-rest encryption and security is available.

See the *NetBackup Security and Encryption Guide*:

<http://www.veritas.com/docs/DOC5332>

<http://www.veritas.com/docs/DOC5332>

About key management for encryption of NetBackup cloud storage

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

NetBackup uses KMS to manage the encryption keys for cloud storage.

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

The following table describes the keys that are required for the 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. 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. 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	<p>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:</p> <p>For cloud storage, the following is the format:</p> <pre>storage_server_name:volume_name</pre> <p>The following items describe the requirements for the key group name components for cloud storage:</p> <ul style="list-style-type: none"> ■ <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. <p>The Disk Pool Configuration Wizard conforms to this format when it creates a key group.</p>
Key record	<p>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.</p> <p>A name for the key record is optional. If you use a key name, you can use any name. Veritas recommends 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.</p>

More information about KMS is available in the *NetBackup Security and Encryption Guide*:

<http://www.veritas.com/docs/DOC5332>

<http://www.veritas.com/docs/DOC5332>

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 71.

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

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

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

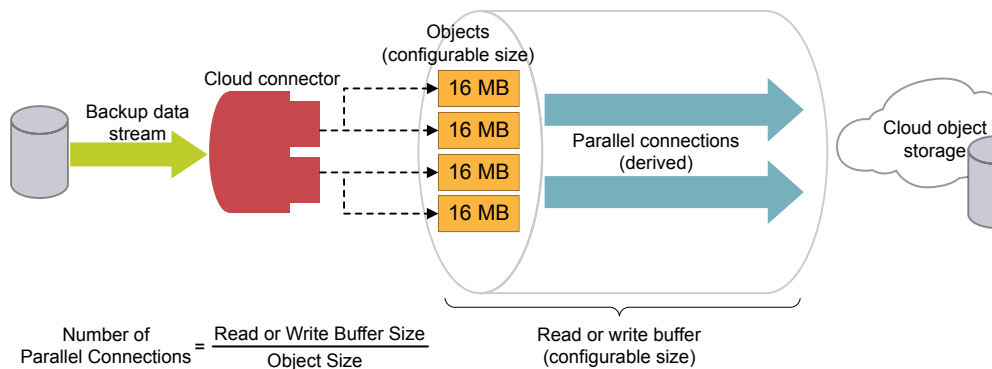
See [“About the NetBackup media servers for cloud storage”](#) on page 97.

About object size for cloud storage

The performance of NetBackup in cloud is driven by the combination of object size, number of parallel connections, and the read or write buffer size. The following diagram illustrates how these factors are related:

The following diagram illustrates how these factors are related:

Figure 3-3 Object size



The parameters are described as follows:

- **Object Size:** The backup data stream is divided into fixed size chunks. These chunks are stored as objects in the cloud object storage. The backup related metadata gets written in variable sizes.
- **Read or write buffer size:** You can configure the read or write buffer size to tune the performance of the backup and restore operations.

Note: 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.

- **Parallel connections (Derived):** To enhance the performance of backup and restore operations, NetBackup uses multiple parallel connections into the cloud storage. The performance of NetBackup depends on the number of parallel connections.

Number of parallel connections is derived from the read or write buffer size and the object size.

Number of Parallel Connections = Read or Write Buffer Size / Object Size

Consider the following factors when deciding the number of parallel connections:

- Maximum number of parallel connections permitted by the cloud storage provider.
- Network bandwidth availability between NetBackup and the cloud storage environment.
- System memory availability on the NetBackup host.

Current default settings

The default settings are as follows:

Table 3-9 Current default settings

Cloud storage provider	CloudCatalyst storage		Classic Cloud storage	
	Object size	Default read or write buffer size	Object size	Default read or write buffer size
Amazon S3/Amazon GovCloud	64 MB (fixed)	64 MB (fixed)	16 MB (fixed)	400 MB (configurable between 16 MB to 1 GB)
Azure	64 MB (fixed)	64 MB (fixed)	4 MB (fixed)	400 MB (configurable between 4 MB to 1 GB)

Considerations

In case of temporary failures on network with data transfer, NetBackup performs multiple retries for transferring the failed objects. In such case, 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 packet loss issues, increasing the number of parallel connections can be helpful.

NetBackup has some timeouts 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.

Consider the following for legacy environments without deduplication support: 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 utilized while also consuming memory. With the increased object size, there is a restriction on number of connections due to the available memory.

If the number of connections are less, parallel downloads would be less compared to older number of connections

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 100.

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 131.

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 132.

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

See “Using media server as NetBackup Cloud master 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:
<http://www.netbackup.com/compatibility>
<http://www.netbackup.com/compatibility>
- 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 master server.

Using media server as NetBackup Cloud master 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:

<http://www.netbackup.com/compatibility>

<http://www.netbackup.com/compatibility>

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

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

To use media server as NetBackup cloud master host

- 1 Identify one of the NetBackup cloud media servers as a cloud master host.
 Choose a media server that has same NetBackup master server version. Do not use a media server with different version.

Note: The media server does not hold the master 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 master host:

```
nbcssc -t -a NetBackup

nbcssc -s -a NetBackup -m cloud_master_host -f
http://www.veritas.com/docs/DOC5332
```

For information on the command, see *Veritas NetBackup Commands Reference Guide*.

- 3 Ensure that the values of **CSSC_PORT** and **CSSC_IS_SECURE** as mentioned in `cloudstore.conf` file from cloud master 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 master host, do not change the name again to point to another media server. If you need to do so, contact Veritas Technical Support.

Additional task post disaster recovery

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

- To perform the task using the NetBackup Administrators Console, see 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
-pxauth_type proxy_auth_type -pxtunnel proxytunnel_usage
http://www.veritas.com/docs/DOC5332
```

For information on the command, see *Veritas NetBackup Commands Reference Guide*.

Additional task post master server upgrade

This is applicable for a NetBackup environment where a master 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 Master host.

After upgrading the master 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://www.veritas.com/support/en_US/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. The NetBackup **Cloud Storage Server Configuration Wizard** communicates with your cloud storage vendor's service endpoint and selects the appropriate host for the storage server.

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

The wizard also lets you enable encryption and configure corresponding parameters for the NetBackup Key Management Service.

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

If you configure encryption, Veritas recommends 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 129.

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* for a complete description about the commands. The guide is available at the following location:

<http://www.veritas.com/docs/DOC5332>

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

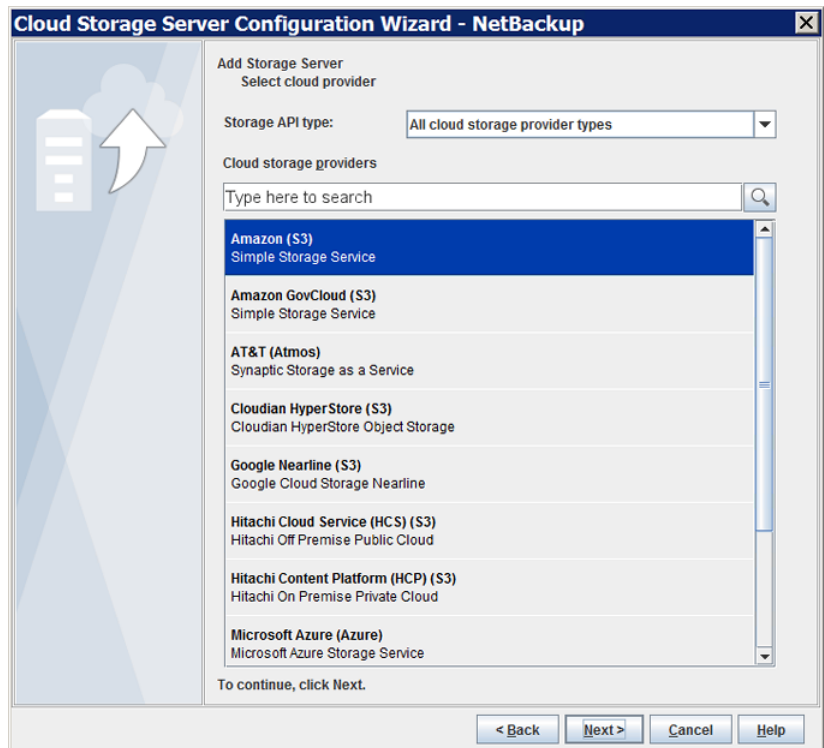
See [“About the NetBackup media servers for cloud storage”](#) on page 97.

To configure a cloud storage server by using the wizard

- 1 In the **NetBackup Administration Console** connected to the NetBackup master server, select either **NetBackup Management** or **Media and Device Management**.
- 2 In the right pane, click **Configure Cloud Storage Servers**.
- 3 Click **Next** on the welcome panel.

The **Select cloud provider** panel appears.

The following is an example of the panel:



- 4 On the **Select cloud provider** panel, perform one of the following:
 - Select the cloud provider from the **Cloud storage providers** list of cloud providers.
 - Sort the list of cloud providers by selecting the cloud storage API type from the **Storage API type** drop-down list and then selecting the cloud provider.

- In the **Cloud storage providers** search box, type the cloud provider name that you want to select. A cloud provider may support multiple cloud storage API types. Select an appropriate provider.
- 5 Click **Next**. A wizard panel for the selected cloud provider appears.
 - 6 On the wizard panel for your cloud provider, select or enter the appropriate information.

The information that is required depends on the cloud vendor. Descriptions of the information that is required for each provider is provided in other topics, based on the storage type API. Those topics also include examples of the wizard panels.

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

See “[About Microsoft Azure cloud storage API type](#)” on page 50.

See “[About OpenStack Swift cloud storage API type](#)” on page 57.

Amazon S3 API <http://www.veritas.com/docs/000115794>

Microsoft Azure API <http://www.veritas.com/docs/000115796>

OpenStack Swift <http://www.veritas.com/docs/000115797>

Note: The provider information topics may include notes, caveats, or warnings. Ensure that you review the topics before you complete the fields in the wizard panel.

- 7 Select the preferred storage class and click **Next**.

Note: This option is available only for Amazon and Amazon GovCloud cloud providers. See “[About Amazon S3 storage classes](#)” on page 31.

- 8 Specify the following settings on the **Specify object size, compression, and encryption settings** panel.

Note: NetBackup media servers that are older than the 7.7.3 version do not support data compression. Therefore, if you have selected an older media server, the compression option does not appear on the panel.

Caution: If you use NetBackup commands to add a NetBackup 7.7.3 or earlier media server to a cloud storage environment that uses compression, cloud backups may fail. Ensure that all media servers that you add to a cloud storage configuration with the compression are NetBackup 7.7.3 or later.

- To specify a custom object size, enter a value in the **Object Size** field. If you do not update the value, the default object size is used.

Note: The object size must be less than or equal to the read or write buffer size.

- To compress your backup data, select **Compress data before writing to cloud storage**.
See [“About data compression for cloud backups”](#) on page 91.
- To encrypt the data that would go on cloud storage, select **Encrypt data using AES-256 before writing to cloud storage**. Then, enter the information to protect the KMS database.
See [“KMS database encryption settings”](#) on page 104.

Click **Next**. If you entered the compression and the encryption information, a dialog box appears that explains that you cannot change the settings after configuration. Click **Yes** to proceed or click **No** to cancel. If you click **Yes**, the **Cloud Storage Server Configuration Summary** panel appears.

- 9 On the **Cloud Storage Server Configuration Summary** panel, verify the selections.

If you need to make corrections, click **Back** until you reach the panel on which you need to make corrections.

If the selections are OK, click **Next**. The wizard creates the storage server, and the **Storage Server Creation Confirmation** panel appears.

- 10 On the **Storage Server Creation Confirmation** panel, do one of the following:
 - To continue to the **Disk Pool Configuration Wizard**, click **Next**.
See [“Configuring a disk pool for cloud storage”](#) on page 120.
 - To exit from the wizard, click **Finish**.
If you exit, you can still create a disk pool.
See [“Configuring a disk pool for cloud storage”](#) on page 120.

KMS database encryption settings

Table 3-10 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. The **Cloud Storage Server Configuration Wizard** and the **Disk Pool Configuration Wizard** configures the encryption for you.

Table 3-10 Required information for the encryption database

Field Name	Required information
KMS Server Name	This field displays the name of your NetBackup master server. You can only configure KMS on your master 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.
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, Veritas recommends 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 129.

See [“Configuring a storage server for cloud storage”](#) on page 100.

Assigning a storage class to Amazon cloud storage

In NetBackup, you can assign a storage class to cloud storage while you configure a new storage server.

See [“About Amazon S3 storage classes”](#) on page 31.

See [“Configuring a storage server for cloud storage”](#) on page 100.

To assign a storage class

- 1 In the NetBackup **Administration Console** > **Cloud Storage Configuration** wizard, select **Amazon**.
- 2 On the **Add Storage Server** screen, specify the Amazon S3 configuration details such as, service host, storage server name, and access details.
- 3 Select the preferred storage class and click **Next**. Veritas recommends that you do not modify the storage class of a cloud storage server after you have assigned it.

See [“About Amazon S3 storage classes”](#) on page 31.

Note: Prior to NetBackup 8.1.1, in the **Advanced Server Configuration** screen, the **x-amz-storage-class** header displayed the Amazon S3 storage classes that NetBackup supports.

Note: **AMZ:STORAGE_CLASS** lists the storage class in the storage server properties dialog box.

- 4 Configure a new disk pool.

See [“Configuring a disk pool for cloud storage”](#) on page 120.

Note: Veritas recommends that you use different buckets for different storage classes.

- 5 Configure a new storage unit by accessing **NetBackup Administration Console > NetBackup Management > Storage > Storage Units**.
- 6 Modify the existing policy or SLP (or create new policy or SLP) to use the new storage unit by accessing the respective user interfaces:
 - To access policy, do the following: In the **NetBackup Administration Console**, expand **NetBackup Management**, and click **Policies**.
 - To access SLP, do the following: In the **NetBackup Administration Console**, expand **NetBackup Management**, expand **Storage**, and click **Storage Life Cycle Policies**.

Changing cloud storage server properties

The Change Storage Server dialog box lists all storage server properties. You can change these properties, if required.

See [“Configuring cloud storage in NetBackup”](#) on page 69.

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

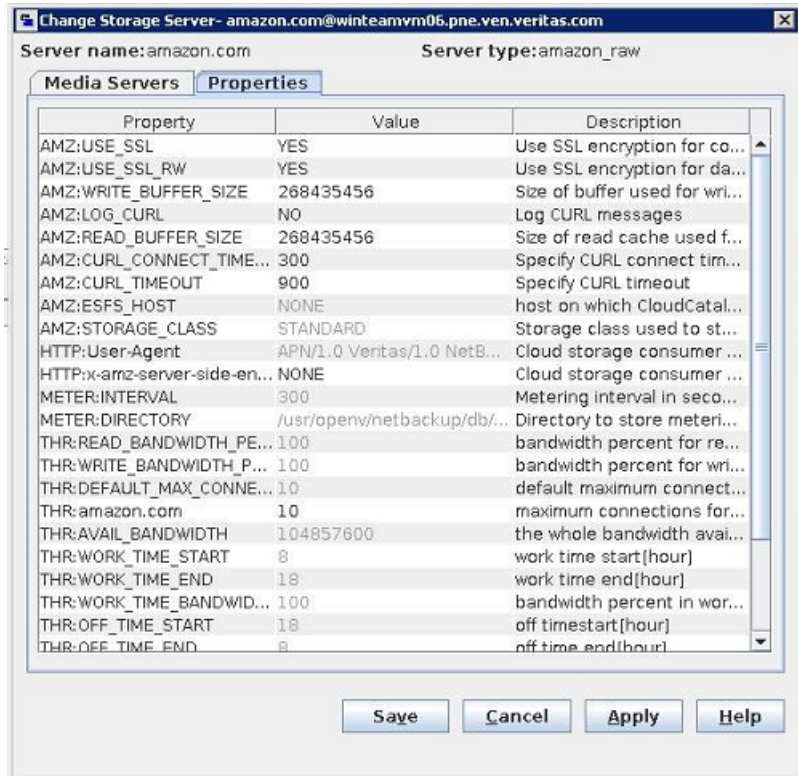
See [“Changing cloud storage host properties”](#) on page 79.

To change cloud storage server properties

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Credentials > Storage Server**.
- 2 Select the storage server.
- 3 On the **Edit** menu, select **Change**.

- 4 In the **Change Storage Server** dialog box, select the **Properties** tab.

The following is an example of the **Properties** for Amazon S3 storage server of type `amazon_raw`:



- 5 To change a property, select its value in the **Value** column and then change it.

See [“NetBackup cloud storage server properties”](#) on page 108.

See [“NetBackup cloud storage server connection properties”](#) on page 112.

See [“NetBackup cloud storage server encryption properties”](#) on page 119.

- 6 Repeat step 5 until you have finished changing the properties.
- 7 Click **OK**.
- 8 Restart the NetBackup Remote Manager and Monitor Service (`nbrmms`) by using the **NetBackup Administration Console Activity Monitor**.

NetBackup cloud storage server properties

The **Properties** tab of the **Change Storage Server** dialog box lets you 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 112.
AMZGOV	Amazon GovCloud	See "NetBackup cloud storage server connection properties" on page 112.
AZR	Microsoft Azure	See "NetBackup cloud storage server connection properties" on page 112.
CLD	Cloudian Hyperstore	See "NetBackup cloud storage server connection properties" on page 112.
CRYPT	Encryption	See "NetBackup cloud storage server encryption properties" on page 119.
GOOG	Google Nearline	See "NetBackup cloud storage server connection properties" on page 112.
HT	Hitachi	See "NetBackup cloud storage server connection properties" on page 112.
HTTP	HTTP headers	See "NetBackup cloud storage server connection properties" on page 112. Note: This field applies to Amazon S3-compatible cloud providers.
METER	Metering	See "NetBackup cloud storage server connection properties" on page 112.
MSDPCLD	CloudCatalyst deduplication to the cloud	See "NetBackup CloudCatalyst storage server properties" on page 118.
ORAC	Oracle Cloud	See "NetBackup cloud storage server connection properties" on page 112.
SWSTK-SWIFT	SwiftStack (Swift)	See "NetBackup cloud storage server connection properties" on page 112.

Table 3-11 Prefix definitions (*continued*)

Prefix	Definition	For more information
THR	Throttling	See “NetBackup cloud storage server bandwidth throttling properties” on page 109.
VER	Verizon	See “NetBackup cloud storage server connection properties” on page 112.

See [“Changing cloud storage server 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.

To change these properties, use the **Scalable Storage** host properties **Cloud Settings** tab.

See [“Scalable Storage properties”](#) on page 71.

Table 3-12 Cloud storage server bandwidth throttling properties

Property	Description
<code>THR:storage_server</code>	<p>Shows maximum number of concurrent jobs that can be run for a specific cloud storage server.</p> <p>If configuring throttling for a media server that is a CloudCatalyst cloud storage server:</p> <ul style="list-style-type: none"> 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 71. <p>Default value: Not applicable</p> <p>Possible values: See the Description column</p>
<code>THR:AVAIL_BANDWIDTH</code>	<p>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.</p> <p>Default value: 104857600</p> <p>Possible values: Any positive integer</p>

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

Property	Description
<p>THR:DEFAULT_MAX_CONNECTIONS</p>	<p>The default maximum number of concurrent jobs that the media server can run for the cloud storage server.</p> <p>If THR:storage_server is set, NetBackup uses THR:storage_server instead of THR:DEFAULT_MAX_CONNECTIONS.</p> <p>This is a read-only field.</p> <p>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.</p> <p>If NetBackup is configured to allow more jobs than THR:DEFAULT_MAX_CONNECTIONS, NetBackup fails any jobs that start after the number of maximum jobs is reached. Jobs include both backup and restore jobs.</p> <p>You can configure job limits per backup policy and per storage unit.</p> <p>See the <i>NetBackup Administrator's Guide, Volume I</i>: http://www.veritas.com/docs/DOC5332</p> <p>Note: NetBackup must account for many factors when it starts jobs: the number of concurrent jobs, the number of THR:DEFAULT_MAX_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.</p> <p>In practice, you should not need to set this value higher than 100.</p> <p>Default value: 10</p> <p>Possible values: 1 to 2147483647</p>
<p>THR:OFF_TIME_BANDWIDTH_PERCENT</p>	<p>This read-only field displays the bandwidth percent that is used during off time.</p> <p>Default value: 100</p> <p>Possible values: 0 to 100</p>

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

Property	Description
THR:OFF_TIME_END	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
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

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

Property	Description
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
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

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

See [“NetBackup cloud storage server properties”](#) on page 108.

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
<p>METER:DIRECTORY</p>	<p>This read-only field displays the directory in which to store data stream metering information.</p> <p>Default value: UNIX: <i>/usr/opensv/var/global/wmc/cloud</i> or <i>/usr/opensv/netbackup/db/cloud</i> (on media server versions 7.7.x to 8.1.2 only)</p> <p>Windows: <i>install_path\Veritas\NetBackup\var\global\wmc\cloud</i> or <i>install_path\Veritas\NetBackup\db\cloud\</i> (on media server versions 7.7.x to 8.1.2 only)</p>
<p>METER:INTERVAL</p>	<p>The interval at which NetBackup gathers connection information for reporting purposes.</p> <p>NetBackup OpsCenter uses the information that is collected to create reports. The value is set in seconds. The default setting is 300 seconds (5 minutes). If you set this value to zero, metering is disabled</p> <p>To change this property, use the Cloud Settings tab of the Scalable Storage host properties.</p> <p>See “Scalable Storage properties” on page 71.</p> <p>Default value: 300</p> <p>Possible values: 1 to 10000</p>
<p>PREFIX:CURL_CONNECT_TIMEOUT</p>	<p>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.</p> <p>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.</p> <p>This value cannot be disabled. If an invalid number is entered, the <code>CURL_CONNECT_TIMEOUT</code> returns to the default value of 300.</p> <p>Default value: 300</p> <p>Possible values: 1 to 10000</p>

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

Property	Description
<i>PREFIX:CURL_TIMEOUT</i>	<p>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 timeout, set the value to 0 (zero).</p> <p>Default value: 900</p> <p>Possible values: 1 to 10000</p>
<i>PREFIX:ESFS_HOST</i>	<p>Identifies the host that contains the ESFS cache. The ESFS cache is used by a CloudCatalyst storage server for deduplication to the cloud.</p> <p>This property is set internally and cannot be changed by the user.</p>
<i>PREFIX:LOG_CURL</i>	<p>Determines if cURL activity is logged. The default is NO which means log activity is disabled.</p> <p>Default value: NO</p> <p>Possible values: NO (disabled) and YES (enabled)</p>
<i>PREFIX:READ_BUFFER_SIZE</i>	<p>The size of the buffer to use for read operations. <i>READ_BUFFER_SIZE</i> is specified in bytes.</p> <p>To enable the use of the buffer, set this value to a non-zero number.</p> <p>The <i>READ_BUFFER_SIZE</i> 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 timeouts. When you calculate the required bandwidth, consider the total load of simultaneous backup jobs and restore jobs on multiple media servers.</p>

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

Property	Description
<i>PREFIX:USE_SSL</i>	<p>Determines if Secure Sockets Layer encryption is used for the control APIs. The default value is YES, meaning SSL is enabled.</p> <p>Default value: YES</p> <p>Possible values: YES or NO</p>
<i>PREFIX:USE_SSL_RW</i>	<p>Determines if Secure Sockets Layer encryption is used for read and write operations. The default value is YES, meaning SSL is enabled.</p> <p>Default value: YES</p> <p>Possible values: YES or NO</p>
<i>PREFIX: WRITE_BUFFER_NUM</i>	<p>This parameter is not applicable for Amazon S3-compatible cloud providers.</p> <p>This read-only field displays the total number of write buffers that are used by the plug-in. The WRITE_BUFFER_SIZE value defines the size of the buffer. The value is set to 1 and cannot be changed.</p> <p>Default value: 1</p> <p>Possible values: 1</p>
<i>PREFIX:WRITE_BUFFER_SIZE</i>	<p>The size of the buffer to use for write operations. WRITE_BUFFER_SIZE is specified in bytes.</p> <p>To disable the use of the buffer, set this value to 0 (zero).</p> <p>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 timeouts. When you calculate the required bandwidth, consider the total load of simultaneous backup jobs and restore jobs on multiple media servers.</p>

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

Property	Description
HTTP:User-Agent	<p>This is applicable only for Amazon S3-compatible cloud providers.</p> <p>This property is set internally and cannot be changed by the user.</p>
HTTP:x-amz-server-side-encryption	<p>This is applicable only for the following cloud providers: Amazon S3 and Amazon GovCloud</p> <p>Use this property to enable the server-side encryption of the data that you need to transfer to the cloud storage.</p> <p>AES-256 is a server-side encryption standard.</p> <p>Set this property to NONE to disable the server-side encryption for the cloud provider.</p> <p>Note: You should not enable this property, if you have already enabled the media server-side encryption option while configuring cloud storage server using the NetBackup Administration Console.</p>
AMZ:REGION_NAME	<p>This is applicable only for Amazon GLACIER_VAULT storage class.</p> <p>Displays the region set during configuration of the storage server.</p> <p>This property is set during configuration of the storage server and cannot be changed by the user.</p>
AMZ:UPLOAD_CLASS	<p>This is applicable only for the LIFECYCLE storage class.</p> <p>Use this property to specify the storage class to back up the data.</p> <p>Default value: STANDARD</p> <p>Possible values: STANDARD or STANDARD_IA</p>
AMZ:RETRIEVAL_RETENTION_PERIOD	<p>This is applicable only for Amazon Glacier.</p> <p>Use this property to specify the retrieval retention period in days.</p>

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

Property	Description
AMZ:TRANSITION_TO_STANDARD_IA_AFTER	<p>This is applicable only for the LIFECYCLE storage class.</p> <p>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.</p> <p>If you have set the UPLOAD_CLASS as STANDARD_IA, the TRANSITION_TO_STANDARD_IA_AFTER must be set to NONE.</p>
AMZ:TRANSITION_TO_GLACIER_AFTER	<p>This is applicable only for the LIFECYCLE storage class.</p> <p>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.</p> <p>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.</p> <p>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.</p>
AMZ:STORAGE_CLASS	<p>This is applicable only for the Amazon S3 cloud providers.</p> <p>Displays the storage class used by the cloud storage server.</p> <p>This property is set internally and cannot be changed by the user.</p>

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

See [“NetBackup cloud storage server properties”](#) on page 108.

NetBackup CloudCatalyst storage server properties

The `MSDPCLD` prefix specifies a deduplication storage property in the **Properties** tab of the **Change Storage Server dialog** box. The following table describes the properties.

Table 3-14 CloudCatalyst storage server properties

Property	Description
MSDPCLD:storagepath	Storage Path
MSDPCLD:spalogpath	Storage Pool Log Path
MSDPCLD:dbpath	Database Path
MSDPCLD:required_interface	Required Interface
MSDPCLD:spalogretention	Storage Pool Log Retention
MSDPCLD:verboselevel	Storage Pool Verbose Level (Range 0 - 5)
MSDPCLD:replication_target(s)	Replication Target(s)
MSDPCLD:dedupetocloud	Dedupe To Cloud
MSDPCLD:Storage Pool Raw Size	Storage Pool Raw Size
MSDPCLD:Storage Pool Reserved Space	Storage Pool Reserved Space
MSDPCLD:Storage Pool Size	Storage Pool Size
MSDPCLD:Storage Pool Used Space	Storage Pool Used Space
MSDPCLD:Storage Pool Available Space	Storage Pool Available Space
MSDPCLD:Catalog Logical Size	Catalog Logical Size
MSDPCLD:Catalog files Count	Catalog files Count
MSDPCLD:Deduplication Ratio	Deduplication Ratio

See [“NetBackup cloud storage server properties”](#) on page 108.

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

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 3-15 Encryption cloud storage server properties

Property	Description
<code>CRYPT:KMS_SERVER</code>	<p>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 master server name. You cannot change this value.</p> <p>Default value: The NetBackup master server name</p> <p>Possible values: N/A</p>
<code>CRYPT:KMS_VERSION</code>	<p>This read-only field displays the NetBackup Key Management Service version. You cannot change this value.</p> <p>Default value: 16</p> <p>Possible values: N/A</p>
<code>CRYPT:LOG_VERBOSE</code>	<p>This read-only field displays if logs are enabled for encryption activities. The value is either <code>YES</code> for logging or <code>NO</code> for no logging.</p> <p>Default value: NO</p> <p>Possible values: YES and NO</p>
<code>CRYPT:VERSION</code>	<p>This read-only field displays the encryption version. You cannot change this value.</p> <p>Default value: 13107</p> <p>Possible values: N/A</p>

See [“NetBackup cloud storage server properties”](#) on page 108.

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

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.

See “[Configuring a disk pool for cloud storage](#)” on page 120.

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

See the *NetBackup Administrator's Guide, Volume I*:

<http://www.veritas.com/docs/DOC5332>

Configuring a disk pool for cloud storage


Use the NetBackup **Disk Pool Configuration Wizard** to create a disk pool for cloud storage. If you create encrypted storage, you must enter a pass phrase for each selected volume that uses encryption. The pass phrase creates the encryption key for that volume.

To configure a cloud storage disk pool by using the wizard

- 1 If the **Disk Pool Configuration Wizard** was launched from the **Storage Server Configuration Wizard**, go to step 5.
Otherwise, in the **NetBackup Administration Console**, select either **NetBackup Management** or **Media and Device Management**.
- 2 From the list of wizards in the right pane, click **Configure Disk Pool**.

- 3 On the **Welcome** panel, the types of disk pools that you can configure depend on the types of storage servers that exist in your environment.

The following is an example of the wizard panel:




Read the information on the welcome panel of the wizard. Then, select the appropriate storage server type and click **Next**.

The **Storage Server Selection** panel appears.

- 4 On the **Storage Server Selection** panel, the storage servers that you configured for the selected storage server type appear.

The following is an example of the wizard panel:



Select the storage server for this disk pool.


After you select the cloud storage server, click **Next**. The **Volume Selection** wizard panel appears.

- 5 The **Volume Selection** panel displays the volumes that have been created already under your account within the vendor's cloud storage.

Note: The following properties do not apply to cloud storage disk pools: **Total available space**, **Total raw 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 is an example of the wizard panel:



To add a volume, click **Add New Volume**. A dialog box appears that contains the information that is required for a volume for your cloud vendor. In that dialog box, enter the required information. Use the following link to find the information about the requirements for the volume names.

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

To select a volume, click the check box for the volume. You can select one volume only.

After you select the volume for the disk pool, click **Next**. The behavior of the wizard depends on whether you configured encryption for the storage server, as follows:

No encryption If you selected a volume on a storage destination that does not require encryption, the **Additional Disk Pool Information** panel appears.

Go to the next step, step 6.

Encryption If you selected a volume on a storage destination that requires encryption, a **Settings** dialog box appears in which you must enter an encryption pass phrase. The pass phrase is for the *key group* key for this storage volume and storage server combination.

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

After you enter a pass phrase and then click **OK** in the **Settings** dialog box, the dialog box closes. Click **Next** in the **Volume Selection** wizard panel to continue to the **Additional Disk Pool Information** wizard panel.

Continue to the next step, step 6.

- On the **Additional Disk Pool Information** panel, enter or select the properties for this disk pool.

The following is an example of the wizard panel:

Disk Pool Configuration Wizard

Additional Disk Pool Information
 Provide additional disk pool information.

Storage server type: vendor_type

Disk Pool Size

Total available space: 8192.00 PB
 Total raw size: 8192.00 PB

Disk Pool name:

Comments:

High water mark: %

Low water mark: %

Maximum I/O Streams

i Concurrent read and write jobs affect disk performance.
 Limit I/O streams to prevent disk overload.

Limit I/O streams: per volume

< Back Next > Cancel Help

See “[Cloud storage disk pool properties](#)” on page 142.

After you enter the additional disk pool information, click **Next**. The **Summary** panel appears.

- 7 On the **Summary** panel, verify the selections.


If the summary shows your selections accurately, click **Next**.

Veritas recommends that you save the KMS key group name and the KMS key name. They are required to recover the keys.

See [“Saving a record of the KMS key names for NetBackup cloud storage encryption”](#) on page 129.

- 8 After NetBackup creates the disk pool, a wizard panel describes the successful action.

The following is an example of the wizard panel:



After NetBackup creates the disk pool, you can do the following:

Configure a storage unit Ensure that **Create a storage unit using the disk pool that you have just created** is selected and then click **Next**. The **Storage Unit Creation** wizard panel appears. Continue to the next step.

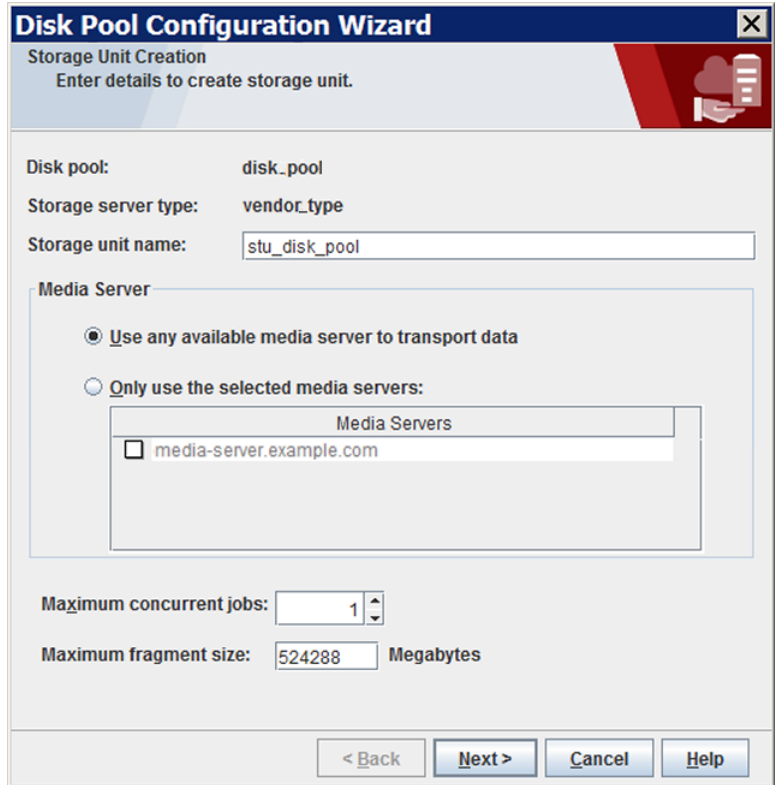
Exit Click **Close**.

You can configure one or more storage units later.

See [“Configuring a storage unit for cloud storage”](#) on page 132.

- 9 On **Storage Unit Creation** wizard panel, enter the appropriate information for the storage unit.

The following is an example of the wizard panel:



See [“Cloud storage unit properties”](#) on page 133.

After you enter or select the information for the storage unit, click **Next** to create the storage unit.

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

See [“Configure a favorable client-to-server ratio”](#) on page 135.

See [“Control backup traffic to the media servers”](#) on page 136.

- 10 After NetBackup configures the storage unit, the **Finished** panel appears. Click **Finish** to exit from the wizard.

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

Veritas recommends 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.

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

To save a record of the key names

- 1** To determine the key group names, use the following command on the master 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
```

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

UNIX: `/usr/opensv/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.

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 In the NetBackup Administration Console, expand **Media and Device Management > Credentials > Storage Servers**.
- 2 Select the cloud storage server.
- 3 From the **Edit** menu, select **Change**.
- 4 In the **Change Storage Server** dialog box, select the **Media Servers** tab.
- 5 Select the media server or servers that you want to enable for cloud backup. The media servers that you select are configured as cloud servers.
- 6 Click **OK**.
- 7 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.

The **Disk Pool Configuration Wizard** lets you create a storage unit; therefore, you may have created a storage unit when you created a disk pool. To determine if storage units exist for the disk pool, see the **NetBackup Management > Storage > Storage Units** window of the Administration Console.

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.

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

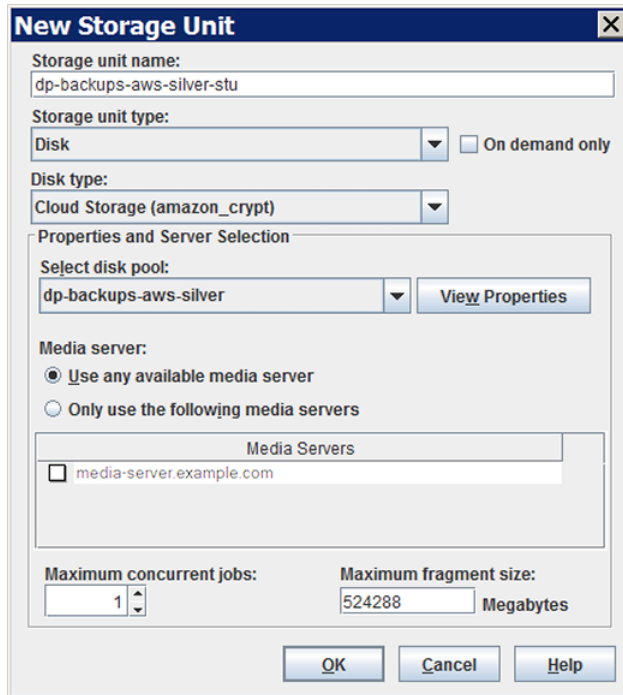
See [“Configure a favorable client-to-server ratio”](#) on page 135.

See [“Control backup traffic to the media servers”](#) on page 136.

See [“Configuring cloud storage in NetBackup”](#) on page 69.

To configure a storage unit from the Actions menu

- 1** In the **NetBackup Administration Console**, expand **NetBackup Management > Storage > Storage Units**.
- 2** On the **Actions** menu, select **New > Storage Unit**.



- 3** Complete the fields in the **New Storage Unit** dialog box.
 See [“Cloud storage unit properties”](#) on page 133.

Cloud storage unit properties

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

Table 3-16 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.

Table 3-16 Cloud storage unit properties (*continued*)

Property	Description
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: <ul style="list-style-type: none"> ■ 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.
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.

Table 3-16 Cloud storage unit properties (*continued*)

Property	Description
Maximum fragment size	<p>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.</p> <p>For a FlashBackup policy, Veritas recommends that you use the default, maximum fragment size to ensure optimal duplication performance.</p>

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.

See “[Configuring a storage unit for cloud storage](#)” on page 132.

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.

See “[Configuring a storage unit for cloud storage](#)” on page 132.

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.

- *Veritas NetBackup Deduplication Guide*
- *Veritas NetBackup Administrator's Guide, Volume I*

These guides are available through the following URL:

<http://www.veritas.com/docs/DOC5332>

<http://www.veritas.com/docs/DOC5332>

Enabling NetBackup Accelerator with cloud storage

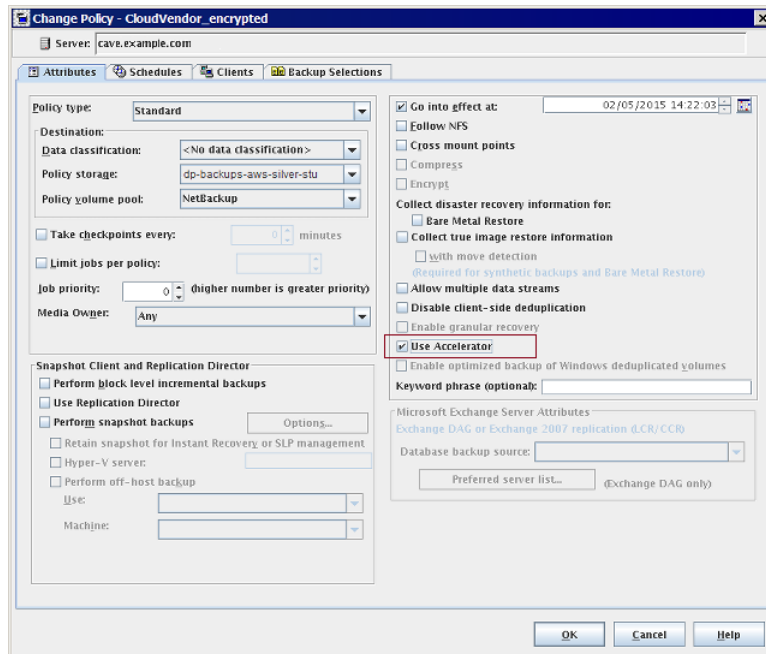
Use the following procedure to enable NetBackup Accelerator for use with NetBackup cloud storage.

Enabling Accelerator for use with NetBackup cloud storage

- 1 In the NetBackup Administration Console, select **NetBackup Management > Policies > *policy_name***. Select **Edit > Change**, and select the **Attributes** tab.
- 2 Select **Use accelerator**.
- 3 Confirm the **Policy storage** option is a valid Cloud storage unit.

The storage unit that is specified under **Policy storage** must be one of the supported Cloud vendors. You can't set **Policy storage** to **Any Available**.

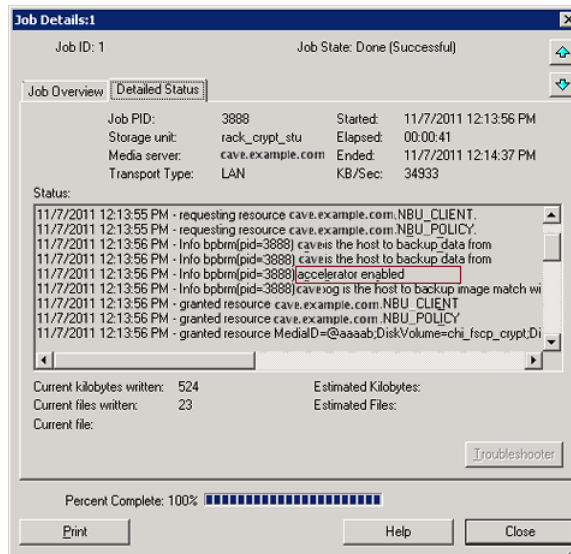
Figure 3-4 Enable Accelerator



Determining if NetBackup Accelerator was used during a backup operation

- 1 In the NetBackup Administration Console, select **Activity Monitor**. Double click the backup that you want to check.
- 2 Click the **Detailed Status** tab.
- 3 Review the status for **accelerator enabled**. This text indicates the backup used NetBackup Accelerator.

Figure 3-5 Confirm Accelerator used during backup



Enabling optimized synthetic backups with cloud storage

Optimized Synthetic backups require three backup schedules. You must have a **Full backup**, an **Incremental backup**, and a **Full Backup with Synthetic backup enabled**. You can use either a Differential incremental or a Cumulative incremental for the incremental backup. You must then perform a full backup, then at least one incremental backup, and finally a full backup with synthetic enabled. The final backup is the optimized synthetic backup.

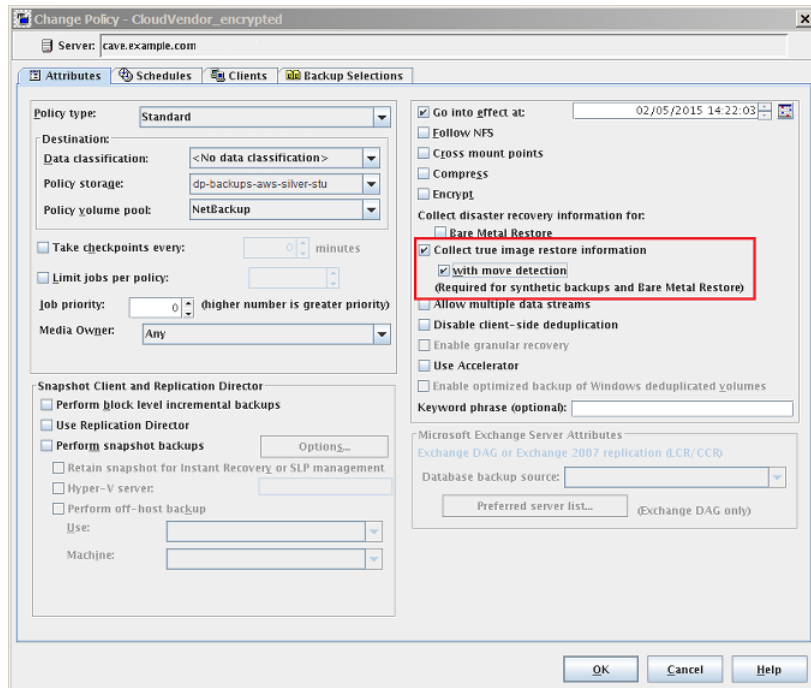
Note: In the case of Hitachi cloud configuration, the True Image Restore (TIR) or synthetic backups do not work, if you have enabled the encryption option. To successfully run the TIR or synthetic backups, you need to enable the versioning option for buckets (or namespaces) through the Hitachi cloud portal. For more details on how to enable the versioning option, contact Hitachi cloud provider.

Enabling Optimized Synthetic backups for use with NetBackup Cloud Storage

- 1 In the NetBackup Administration Console, select **NetBackup Management > Policies > *policy_name***. Select **Edit > Change**, and select the **Attributes** tab.
- 2 Select **Collect true image restore information** and **with move detection**.
- 3 Confirm the **Policy storage** option is a valid Cloud storage unit.

The storage unit that is specified under **Policy storage** must be one of the supported Cloud vendors. You can't set **Policy storage** to **Any Available**.


Figure 3-6 Enable Optimized Synthetic backups



Determining if a backup was an Optimized Synthetic backup

- 1 In the NetBackup Administration Console, select **Activity Monitor**. Double click the backup that you want to check.
- 2 Click the **Detailed Status** tab.
- 3 Review the status for **Performing Optimized Synthetic Operation**. This text indicates the backup was an Optimized Synthetic backup.

Figure 3-7 Confirm backup was Optimized Synthetic



Creating a backup policy

The easiest method to set up a backup policy is to use the **Policy Configuration Wizard**. This wizard guides you through the setup process by automatically choosing the best values for most configurations.

Not all policy configuration options are presented through the wizard. For example, calendar-based scheduling and the **Data Classification** setting. After the policy is created, modify the policy in the **Policies** utility to configure the options that are not part of the wizard.

Note: Do not use the Policy Configuration Wizard to configure policies for Replication Director.

Using the Policy Configuration Wizard to create a backup policy

Use the following procedure to create a backup policy with the Policy Configuration Wizard.

To create a backup policy with the Policy Configuration Wizard

- 1 In the **NetBackup Administration Console**, in the left pane, click **NetBackup Management**.
- 2 In the right pane, click **Create a Policy** to begin the **Policy Configuration Wizard**.
- 3 Select **File systems, databases, applications**.
- 4 Click **Next** to start the wizard and follow the prompts.

Click **Help** on any wizard panel for assistance while running the wizard.

Creating a backup policy without using the Policy Configuration Wizard

Use the following procedure to create a backup policy in the **NetBackup Administration Console** without using the Policy Configuration Wizard.

To create a policy without the Policy Configuration Wizard

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 On the **Actions** menu, click **New > Policy**.
- 3 Type a unique name for the new policy in the **Add a New Policy** dialog box.
- 4 If necessary, clear the **Use Policy Configuration Wizard** check box.
- 5 Click **OK**.
- 6 Configure the attributes, the schedules, the clients, and the backup selections for the new policy.


Changing cloud storage disk pool properties

You can change some of the properties of a disk pool.

To change disk pool properties

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices > Disk Pools**.
- 2 Select the disk pool that you want to change in the details pane.

- 3 On the **Edit** menu, select **Change**.



- 4 Change the properties as necessary.
 See [“Cloud storage disk pool properties”](#) on page 142.
- 5 Click **OK**.

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 3-17 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.
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.

Table 3-17 Cloud storage disk pool properties (*continued*)

Property	Description
Limit I/O streams	<p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p>
per volume	<p>Select or enter the number of read and write streams to allow per volume.</p> <p>Many factors affect the optimal number of streams. Factors include but are not limited to disk speed, CPU speed, and the amount of memory.</p> <p>For the disk pools that are configured for Snapshot and that have a Replication source property:</p> <ul style="list-style-type: none"> ■ 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.

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:**
`install_path\Veritas\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:** `/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.

To add a CA

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

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

Add the following information block:

```
Certificate Authority Name
=====
-----BEGIN CERTIFICATE-----
<Certificate content>
-----END CERTIFICATE-----
```

- 3 Save the file.

To remove a CA

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

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

```
Certificate Authority Name
=====
-----BEGIN CERTIFICATE-----
<Certificate content>
-----END CERTIFICATE-----
```

- 3 Save the file.

List of CAs approved by NetBackup

- Baltimore CyberTrust Root
- Cybertrust Global Root
- DigiCert Assured ID Root CA
- DigiCert Assured ID Root G2
- DigiCert Assured ID Root G3
- DigiCert Global Root CA
- DigiCert Global Root G2
- DigiCert Global Root G3
- DigiCert High Assurance EV Root CA
- DigiCert Trusted Root G4
- GeoTrust Global CA
- GeoTrust Global CA 2
- 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:

- [About monitoring and reporting for cloud backups](#)
- [Viewing cloud storage job details](#)
- [Viewing the compression ratio](#)
- [Viewing NetBackup cloud storage disk reports](#)
- [Displaying KMS key information for cloud storage encryption](#)

About monitoring and reporting for cloud backups

Veritas provides several methods to monitor and report NetBackup cloud storage and cloud storage activity, as follows:

NetBackup OpsCenter The NetBackup OpsCenter provides the most detailed reports of NetBackup cloud storage activity. See the *NetBackup OpsCenter Administrator's Guide* for details on cloud monitoring and reporting:

<http://www.veritas.com/docs/DOC5332>

<http://www.veritas.com/docs/DOC5332>

If OpsCenter cannot connect to the media server, it cannot obtain the necessary data for reporting. Therefore, ensure that the following services are up and running on all the media servers that are configured for cloud storage:

- NetBackup CloudStore Service Container (`nbcssc`) (on media server versions 7.7.x to 8.1.2 only)
- NetBackup Service Layer (`nbsl`) service

Note: Where Amazon is the cloud service provider, OpsCenter cannot report on the data that MSDP cloud storage servers upload to the cloud.

See “[Connection to the NetBackup CloudStore Service Container fails](#)” on page 169.

The NetBackup Administration Console **Disk Pools** window

The **Disk Pools** window displays the values that were stored when NetBackup polled the disk pools. NetBackup polls the disk pools every five minutes.

To display the window, in the **NetBackup Administration Console**, in the left pane, select **Media and Device Management > Devices > Disk Pools**.

Note: The information that is displayed for **Used Capacity** and **Available Space** is inaccurate in the **NetBackup Administration Console**. Even if there is data in the disk pool, the value that is displayed for **Used Capacity** is zero. The value for **Available Space** displays the maximum amount. You must review the information on the provider website for accurate use information.

Note: The information that is displayed for **Used Capacity** and **Available Space** for Amazon is inaccurate in the NetBackup Administration Console. The values are found under **Media and Device Management > Devices > Disk Pool**. Even if there is information in the disk pool, the value that is displayed for **Used Capacity** is zero. The value for **Available Space** displays the maximum amount. You must review the information on the provider website for accurate use information.

NetBackup disk reports See “[Viewing NetBackup cloud storage disk reports](#)” on page 151.

Viewing cloud storage job details

Use the NetBackup Activity Monitor to view job details.

To view cloud storage job details

- 1 In the **NetBackup Administration Console**, click **Activity Monitor**.
- 2 Click the **Jobs** tab.
- 3 To view the details for a specific job, double-click on the job that is displayed in the **Jobs** tab pane.
- 4 In the **Job Details** dialog box, click the **Detailed Status** tab.

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 150.
- 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 bytes to data in bytes bytes,
compression ratio ratio_value
```

```
date:time <PID> <4> 35:bptm:<PID>:
media_server_IP: compress: image image_name_C1_HDR
compressed from data in bytes to data in bytes bytes,
compression ratio ratio_value
```

Viewing NetBackup cloud storage disk reports

The NetBackup disk reports include information about the disk pools, disk storage units, disk logs, and images that are stored on disk media.

[Table 4-1](#) describes the disk reports available.

Table 4-1 Disk reports

Report	Description
Images on Disk	<p>The Images on Disk report generates the image list present on the disk storage units that are connected to the media server. The report is a subset of the Images on Media report; it shows only disk-specific columns.</p> <p>The report provides a summary of the storage unit contents. If a disk becomes bad or if a media server crashes, this report can let you know what data is lost.</p>

Table 4-1 Disk reports (*continued*)

Report	Description
Disk Logs	The Disk Logs report displays the media errors or the informational messages that are recorded in the NetBackup error catalog. The report is a subset of the Media Logs report; it shows only disk-specific columns.
Disk Storage Unit Status	The Disk Storage Unit Status report displays the state of disk storage units in the current NetBackup configuration. Multiple storage units can point to the same disk pool. When the report query is by storage unit, the report counts the capacity of disk pool storage multiple times.
Disk Pool Status	The Disk Pool Status report displays the state of disk pool storage units. This report displays only when a license is installed that enables a NetBackup disk feature.

See [“About monitoring and reporting for cloud backups”](#) on page 148.

To view disk reports

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Reports > Disk Reports**.
- 2 Select the name of a disk report.
- 3 In the right pane, select the report settings.
- 4 Click **Run Report**.

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: Veritas recommends 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/opensv/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        : -
```

Operational notes

This chapter includes the following topics:

- [NetBackup bpstsinfo command operational notes](#)
- [Unable to configure additional media servers](#)
- [Cloud configuration may fail if NetBackup Access Control is enabled](#)
- [Deleting cloud storage server artifacts](#)

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	<p>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.</p> <p>The following is an example of a command that specifies encrypted storage:</p> <pre>bpstsinfo -lsuinfo -storage_server amazon.com -stype amazon_crypt</pre>

Unable to configure additional media servers

If you attempt to run the **Cloud Storage Server Configuration Wizard** on a second media server that uses the same master server as the first media server, the operation fails. An `illegal duplication` error similar to the following appears:



Your only options in the wizard are to click **Cancel** or **Back**. If you click **Back**, there are no configuration changes that allow the wizard to continue.

You must use the correct procedure if you want multiple media servers in your cloud environment. More information is available in a different topic.

See [“To add backup media servers to your cloud environment”](#) on page 132.

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:

<http://www.veritas.com/docs/DOC5332>

<http://www.veritas.com/docs/DOC5332>

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.

Troubleshooting

This chapter includes the following topics:

- [About unified logging](#)
- [About legacy logging](#)
- [NetBackup cloud storage log files](#)
- [Enable libcurl logging](#)
- [NetBackup Administration Console fails to open](#)
- [Troubleshooting cloud storage configuration issues](#)
- [Troubleshooting cloud storage operational issues](#)

About unified logging

Unified logging and legacy logging are the two forms of debug logging used in NetBackup. All NetBackup processes use one of these forms of logging. Server processes and client processes use unified logging.

Unified logging creates log file names and messages in a standardized format. These logging files cannot be easily viewed with a text editor. They are in binary format and some of the information is contained in an associated resource file. Only the `vxlogview` command can assemble and display the log information correctly.

Unlike legacy logging, unified logging does not require that you create logging subdirectories. 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`

You can access logging controls in the **NetBackup Administration Console**. In the left pane, expand **NetBackup Management > Host Properties > Master Servers** or **Media Servers**. Double-click the server you want to change. In the left pane of the dialog box, click **Logging**.

You can also manage unified logging by using the following commands:

- `vxlogcfg` Modifies the unified logging configuration settings.
 for more information about the `vxlogcfg` command.
- `vxlogmgr` Manages the log files that the products that support unified logging generate.
 for more information about the `vxlogmgr` command.
- `vxlogview` Displays the logs that unified logging generates.
 See [“Examples of using vxlogview to view unified logs”](#) on page 160.
 for more information about the `vxlogview` command.

These commands are located in the following directory:

Windows `install_path\NetBackup\bin`
 UNIX `/usr/opensv/netbackup/bin`

See the [NetBackup Commands Reference Guide](#) for a complete description about these commands.

More information about legacy logging is available.

See [“About using the vxlogview command to view unified logs”](#) on page 159.

About using the vxlogview command to view unified logs

Use the `vxlogview` command to view the logs that unified logging creates. These logs are stored in the following directory.

UNIX `/usr/opensv/logs`
 Windows `install_path\NetBackup\logs`

Unlike the files that are written in legacy logging, unified logging files cannot be easily viewed with a text editor. The unified logging files are in binary format, and

some of the information is contained in an associated resource file. Only the `vxlogview` command can assemble and display the log information correctly.

You can use `vxlogview` to view NetBackup log files as well as PBX log files.

To view PBX logs using the `vxlogview` command, do the following:

- Ensure that you are an authorized user. For UNIX and Linux, you must have root privileges. For Windows, you must have administrator privileges.
- To specify the PBX product ID, enter `-p 50936` as a parameter on the `vxlogview` command line.

`vxlogview` searches all the files, which can be a slow process. Refer to the following topic for an example of how to display results faster by restricting the search to the files of a specific process.

See [“Examples of using vxlogview to view unified logs”](#) on page 160.

Examples of using vxlogview to view unified logs

The following examples demonstrate how to use the `vxlogview` command to view unified logs.

Table 6-1 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 (<code>nbpem</code>) 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>
Display the log messages from a specific time period	Display the log messages for <code>nbpem</code> that were issued during the specified time period: <code># vxlogview -o nbpem -b "05/03/15 06:51:48 AM" -e "05/03/15 06:52:48 AM"</code>

Table 6-1 Example uses of the vxlogview command (*continued*)

Item	Example
Display results faster	<p>You can use the <code>-i</code> option to specify an originator for a process:</p> <pre># vxlogview -i nbpem</pre> <p>The <code>vxlogview -i</code> option searches only the log files that the specified process (<code>nbpem</code>) 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.</p> <p>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.</p> <p>The <code>-i</code> option displays entries for all OIDs that are part of that process including libraries (137, 156, 309, etc.).</p>
Search for a job ID	<p>You can search the logs for a particular job ID:</p> <pre># vxlogview -i nbpem grep "jobid=job_ID"</pre> <p>The <code>jobid=</code> search key should contain no spaces and must be lowercase.</p> <p>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 (<code>nbpem</code>). 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>.</p>

See the *NetBackup Commands Reference Guide* for a complete description of the `vxlogview` command. The guide is available through the following URL:

<http://www.veritas.com/docs/DOC5332>

See “[About using the vxlogview command to view unified logs](#)” on page 159.

About legacy logging

Legacy logging and unified logging are the two forms of debug logging used in NetBackup. All NetBackup processes use either unified logging or legacy logging.

See “[About unified logging](#)” on page 158.

In legacy debug logging, each process creates log files of debug activity in its own logging directory. The NetBackup legacy debug log directories are located in the following directories:

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>

These top-level directories can contain a directory for each NetBackup process that uses legacy logging. By default, NetBackup creates only a subset of all of the possible log directories. For example, the following directories are created by default on UNIX servers:

- nbfp
- nbliveup
- nblogadm
- user_ops

To enable logging for all of the NetBackup processes that use legacy logging, you must create the log file directories that do not already exist, unless you use the Logging Assistant. For more information about the Logging Assistant, see the *NetBackup Administrator's Guide, Volume I*. The guide is available at the following location:

<http://www.veritas.com/docs/DOC5332>

You can use the following batch files to create all of the debug log directories at once:

- Windows: `install_path\NetBackup\Logs\mklogdir.bat`
- UNIX: `/usr/opensv/netbackup/logs/mklogdir`

See the *NetBackup Commands Reference Guide* for a complete description about the `mklogdir` command. The guide is available at the following location:

<http://www.veritas.com/docs/DOC5332>

After the directories are created, NetBackup creates log files in the directory that is associated with each process. A debug 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.

To enable debug logging for the NetBackup Status Collection Daemon (`vmscd`), create the following directory before you start `nbemm`.

Windows	<code>install_path\Volmgr\debug\vmscd\</code>
UNIX	<code>/usr/opensv/volmgr/debug/vmscd</code>

As an alternative, you can restart `vmscd` after creating the directory.

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 master 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*, available through the following URL:

<http://www.veritas.com/docs/DOC5332>

<http://www.veritas.com/docs/DOC5332>

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 Veritas 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 log (VxUL) files. 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*:

<http://www.veritas.com/docs/DOC5332>

<http://www.veritas.com/docs/DOC5332>

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 Veritas representative. A log level of 5 is for troubleshooting only.

Specify the NetBackup log levels in the **Logging** host properties on the NetBackup master server. The log levels for some processes specific to certain options are set in configuration files as described in [Table 6-2](#).

[Table 6-2](#) describes the logs.

Table 6-2 NetBackup logs for cloud storage

Activity	OID	Processes
Backups and restores	N/A	<p>Messages appear in the log files for the following processes:</p> <ul style="list-style-type: none"> ■ 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. <p>The log files reside in the following directories:</p> <ul style="list-style-type: none"> ■ 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	<p>The <code>bpdbm</code> database manager log files.</p> <p>The log files reside in the following directories:</p> <ul style="list-style-type: none"> ■ 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	<p>This is applicable to media server versions 7.7.x to 8.1.2 only.</p> <p>The NetBackup Cloud Storage Service Container (<code>nbcssc</code>) writes log files to the following directories:</p> <ul style="list-style-type: none"> ■ For Windows: <code>install_path\Veritas\NetBackup\logs\nbcssc</code> ■ For UNIX: <code>/usr/opensv/netbackup/logs/nbcssc</code>
NetBackup Web Management Console	495	<p>The NetBackup Web Management Console (<code>nbwmc</code>) service writes logs to the following directories:</p> <ul style="list-style-type: none"> ■ For Windows: <code>install_path\Veritas\netbackup\logs\nbwebsevice</code> ■ For UNIX: <code>/usr/opensv/logs/nbwebsevice</code>

Table 6-2 NetBackup logs for cloud storage (*continued*)

Activity	OID	Processes
NetBackup Service Layer	N/A	The NetBackup Service Layer (<code>nbsl</code>) service writes logs to the following directories: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> For Windows: <code>install_path\Veritas\netbackup\logs\bcssc</code> For UNIX: <code>/usr/openv/netbackup/logs/bcssc</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 175.

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`.

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

NetBackup Administration Console fails to open

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

If you change the port number used by the NetBackup CloudStore Service Container (`nbcssc`), the **NetBackup Administration Console** may not open.

You must change the port number value to 5637 in the following places:

The CloudStore Service Container configuration file

The CloudStore Service Container configuration file resides in the following directories:

- UNIX: `/usr/opensv/java/cloudstorejava.conf`
- Windows:
`install_path\Veritas\NetBackup\bin\cloudstorewin.conf`

The port number is defined in the configuration file as follows:

```
[NBCSSC]
NBCSSC_PORT=5637
```

Note: Port 5637 is used to provide back-level media support for media servers that are configured for cloud storage. Ensure that you make the port number change at all places. Communication with the master server fails if the older media servers use a different port.

The operating system's `services` file

The `services` file is in the following locations:

- Windows:
`C:\WINDOWS\system32\drivers\etc\services`
- Linux: `/etc/services`

For a media server that is promoted as a cloud master, make sure that the port number is the same across all places. If you change the value in the CloudStore Service Container configuration file, ensure that you also change the value in the `services` file.

See [“Connection to the NetBackup CloudStore Service Container fails”](#) on page 169.

Troubleshooting cloud storage configuration issues

The following sections may help you troubleshoot configuration issues.

See [“NetBackup Scalable Storage host properties unavailable”](#) on page 168.

See [“Connection to the NetBackup CloudStore Service Container fails”](#) on page 169.

See [“Cannot create a cloud storage disk pool”](#) on page 170.

See [“Cannot create a cloud storage”](#) on page 171.

See [“NetBackup Administration Console fails to open”](#) on page 167.

See [“Data transfer to cloud storage server fails in the SSL mode”](#) on page 172.

See [“Amazon GovCloud cloud storage configuration fails in non-SSL mode”](#) on page 172.

See [“Data restore from the Google Nearline storage class may fail”](#) on page 173.

See [“Fetching storage regions fails with authentication version V2”](#) on page 174.

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 180.

See [“Stopping and starting the NetBackup CloudStore Service Container”](#) on page 179.

Connection to the NetBackup CloudStore Service Container fails

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

The NetBackup cloud storage `csconfig` configuration command makes three attempts to connect to the NetBackup CloudStore Service Container with a 60-second time-out for each connection attempt. The NetBackup OpsCenter also connects to the NetBackup CloudStore Service Container to obtain data for reporting.

If they cannot establish a connection, verify the following information:

- The NetBackup CloudStore Service Container is active.
 See [“NetBackup CloudStore Service Container startup and shutdown troubleshooting”](#) on page 180.
 See [“Stopping and starting the NetBackup CloudStore Service Container”](#) on page 179.
- Your firewall settings are appropriate.
- The **Enable insecure communication with 8.0 and earlier hosts** option on the NetBackup master server is selected if the media server is of the version 8.0 or earlier. The option is available in the **NetBackup Administration Console** on the **Security Management > Global Security Settings > Secure Communication** tab.
- The `cacert.pem` file is present on both NetBackup master and media server in following locations:

- UNIX/Linux - `/usr/opensv/var/webtruststore`
- Windows - `<install_path>/var/webtruststore`

If the `cacert.pem` file is not present on the master server or a media server, run the `nbcertcmd -getCACertificate` command on that host. After running this command, restart the NetBackup CloudStore Service Container on that host. See the *NetBackup Commands Reference Guide* for a complete description of the command.

Note: This `cacert.pem` file contains the CA certificates that the NetBackup authorization service generates.

- The `cacert.pem` file is same on the NetBackup master and media server.
- The security 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 master server. After running this command, restart the NetBackup CloudStore Service Container on the master server and the media servers.

See “[Deploying host name-based certificates](#)” on page 88.

- If the master server runs on an operating system that does not support NetBackup cloud configurations: You can choose to use the NetBackup CloudStore Service Container on a media server as the master service container. To do so, update the `CSSC_MASTER_NAME` parameter of the `cloudstore.conf` file on all the cloud-supported media servers with the media server name you chose earlier. However, communication from other media servers to the media server that now functions as the master configuration for the `nbcssc` service and vice versa fails. The failure happens because both these media servers verify if a trusted host has made the communication request.

Note: The media server that now functions as the master configuration for the `nbcssc` service must run the same NetBackup version as the NetBackupmaster server.

For the operating systems that NetBackup supports for cloud storage, see the NetBackupoperating system compatibility list available through the following URL:

<http://www.netbackup.com/compatibility>

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

To fix this issue, add the authorized host entries on the media and the master servers that support cloud configurations.

See the 'Adding a server to a servers list' topic in the *NetBackup™ Administrator's Guide, Volume I* for detailed steps.

- 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.

Cannot create a cloud storage disk pool

The following table describes potential solutions if you cannot create a disk pool in NetBackup.

Table 6-3 Cannot create disk pool solutions

Error	Description
<p>The wizard is not able to obtain Storage Server information. Cannot connect on socket. (25)</p>	<p>The error message appears in the Disk Configuration Wizard.</p> <p>The Disk Configuration Wizard query to the cloud vendor host timed-out. The network may be slow or a large number of objects (for example, buckets on Amazon S3) may exist.</p> <p>To resolve the issue, use the NetBackup <code>nbdevconfig</code> command to configure the disk pool. Unlike the wizard, the <code>nbdevconfig</code> command does not monitor the command response times.</p> <p>See the <i>NetBackup Commands Reference Guide</i> for a complete description of the commands. The guide is available at the following location: http://www.veritas.com/docs/DOC5332</p>

Cannot create a cloud storage

If you cannot create a cloud storage in NetBackup, verify the following:

- The `cacert.pem` file is present on both NetBackup master 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 `cacert.pem` file is not present, run the `nbcertcmd -getCACertificate` on the master 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 `cacert.pem` file is a NetBackup-specific file. This file includes the CA certificates generated by the NetBackup authorization service.

- The `cacert.pem` file is same on the NetBackup master 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 master server. After running this command, restart the NetBackup CloudStore Service Container on the master and media server.

See “[Deploying host name-based certificates](#)” on page 88.

- 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 179.
- The **Enable insecure communication with 8.0 and earlier hosts** option on the NetBackup master server is selected if the media server is of the version 8.0 or earlier. The option is available in the **NetBackup Administration Console** on the **Security Management > Global Security Settings > 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.

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. The guide is available at the following location:

<http://www.veritas.com/docs/DOC5332>

<http://www.veritas.com/docs/DOC5332>

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 “[Changing cloud storage server properties](#)” on page 106.

See “[NetBackup cloud storage server connection properties](#)” on page 112.

Backups may fail for cloud storage configurations with Frankfurt region

NetBackup 7.7.1 and later versions support configuring cloud storage using the Frankfurt region. NetBackup media servers that are older than the 7.7.1 version do not support configuring cloud storage using the Frankfurt region.

Cloud backups may fail in the following scenario:

You have configured cloud storage server with a media server that is older than NetBackup 7.7.1. You have created a disk pool in the Frankfurt region using an existing bucket.

To avoid such cloud backup failures, ensure that when you configure cloud storage using the Frankfurt region, the cloud media server is NetBackup 7.7.1 or later version.

Backups may fail for cloud storage configurations with the cloud compression option

The NetBackup cloud data compression option requires all cloud media servers that are associated with the cloud storage configuration to be NetBackup 7.7.3 or later version.

Cloud backups may fail in the following cloud compression scenario:

You have configured cloud storage server using the **NetBackup Administration Console** or the command-line interface with the compression option enabled, with a media server that is compatible. You then add a media server of a version that is older than NetBackup 7.7.3 using the command-line interface, to the same cloud configuration.

To avoid such cloud backup failures, ensure that all media servers that you add to the cloud storage configuration with the compression option to be NetBackup 7.7.3 or later version.

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).

Troubleshooting cloud storage operational issues

The following sections may help you troubleshoot operational issues.

See [“NetBackup Scalable Storage host properties unavailable”](#) on page 168.

See [“Cloud storage backups fail”](#) on page 175.

See [“A restart of the nbcssc \(on legacy media servers\), nbwmc, and nbsl processes reverts all cloudstore.conf settings”](#) on page 180.

See [“NetBackup CloudStore Service Container startup and shutdown troubleshooting”](#) on page 180.

See [“NetBackup Administration Console fails to open”](#) on page 167.

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. E
rrno = 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 Veritas 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.

You must use the NetBackup **Disk Pool Configuration Wizard** to create the volume on the cloud storage. The wizard 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 by using the **Disk Pool Configuration Wizard**.

See ["Viewing cloud storage job details"](#) on page 150.

See [“NetBackup cloud storage log files”](#) on page 163.

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 180.

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 master 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 master 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 master server from the **Media Servers** pane.

See [“Troubleshooting cloud storage operational issues”](#) on page 175.

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 `cacert.pem` file is present on both NetBackup master 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 `cacert.pem` file is not present, run the `nbcertcmd -getCACertificate` on the master 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 `cacert.pem` file is a NetBackup-specific file. This file includes the CA certificates generated by the NetBackup authorization service.

- The `cacert.pem` file is same on the NetBackup master 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 `bpbpbaz -ProvisionCert` on the master server. After running this command, restart the NetBackup CloudStore Service Container on the master and media server.

See [“Deploying host name-based certificates”](#) on page 88.

- For media server versions 7.7.x to 8.1.2, the NetBackup CloudStore Service is active.
See [“Stopping and starting the NetBackup CloudStore Service Container”](#) on page 179.
- The **Enable insecure communication with 8.0 and earlier hosts** option on the NetBackup master server is selected if the media server is of the version 8.0 or earlier. The option is available in the **NetBackup Administration Console**

on the **Security Management > Global Security Settings > 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 112.
- Reduce the number of read/write buffers.
 See [“NetBackup cloud storage server bandwidth throttling properties”](#) on page 109.
- 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.

Use the **NetBackup Administration Console** to 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 180.

To start or stop the CloudStore Service Container

- 1 In the **NetBackup Administration Console**, expand **NetBackup Administration > Activity Monitor**.
- 2 Click the **Daemons** tab (UNIX or the **Services** tab (Windows)).
- 3 In the **Details** pane, select **nbcssc** (UNIX and Linux) or **NetBackup CloudStore Service Container** (Windows).
- 4 On the **Actions** menu, select **Stop Selected** or **Start Selected** (Windows) or **Stop Daemon** or **Start Daemon** (UNIX).

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](#)
- [\[Unresolved xref\]](#)

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 master 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 84.

See [“Stopping and starting the NetBackup CloudStore Service Container”](#) on page 179.

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://isearch.veritas.com/internal-search/en_US/article.100042245.html.

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://isearch.veritas.com/internal-search/en_US/article.100042314.html.

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 check point 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 that 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 Administrator console, 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.

The following error is displayed:



This error occurs while creating a storage server. If you are using the CLI, `tpcommand` to add credential fails.

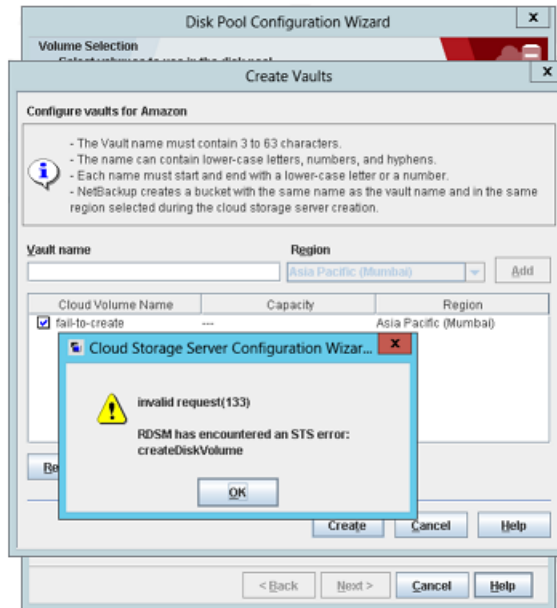
Check the `tpcommand` logs for `AccessDeniedException`, for example,

```
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> magmavm1.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> magmavm1.abc.xyz.qwe.com:
```

- Create Vault or Create Bucket permission (glacier:CreateVault or glacier:DescribeVault) missing.

The following error is displayed:




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.

The following error is displayed:




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. The following error is displayed:



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.

The following error is displayed:




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. The following error is displayed:



This missing permission causes the image cleanup or image expiry process to fail.

Check the bpdm 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.

Index

A

- Add at least one index marker 68
- Amazon
 - glacier vault 42–43
- amazon
 - virtual private cloud 32
- Amazon GLACIER
 - long-term retention 35
- Amazon lifecycle
 - restore 48
- Amazon S3
 - about 15
 - configuration options 19
 - configuration options (advanced) 25
 - credential broker details 29
 - requirements 16

B

- backups fail
 - Accelerator backups fail 175
 - after the WRITE_BUFFER_SIZE is increased 176
 - storage volume was created by the cloud vendor interface 176
 - The NetBackup CloudStore Service Container is not active 177
 - Use any available media server option 177
- bandwidth
 - throttling 109
- bpstsinfo command
 - operational notes 155

C

- catalog
 - cloud configuration files 12
- Certificate Authority (CA) 90
- cloud
 - storage unit properties 133
- cloud configuration files 12

- cloud disk pool
 - changing properties 141
- cloud master host 98
- Cloud Settings tab 71
- cloud storage
 - Amazon S3 API type 15
 - configuring 69
 - Microsoft Azure API type 50
 - OpenStack Swift API type 57
- Cloud Storage host properties 76
- cloud storage instance
 - add 78
 - change 79
 - delete 80
 - manage 79
 - remove 79
- cloud storage properties
 - change 79
 - manage 79
 - remove 79
- cloud storage server
 - about 95
 - bandwidth properties 109
 - changing properties 106
 - CloudCatalyst 118
 - connection properties 112
 - encryption properties 119
 - properties 108
- CloudCatalyst
 - configuring throttling for 109
 - ESFS_HOST cloud connection property 114
 - Maximum concurrent jobs Scalable Storage property 73
- CloudStore Service Container
 - about 81
 - configuring port number 85
 - port number 81
 - security certificate for 82
 - security mode changed while service is active 181
 - security modes 84
 - startup and shutdown troubleshooting 180

cloudstore.conf configuration file 84

Configuration

Accelerator 137

configuration

disk pool configuration wizard 120

optimized synthetic backups for cloud storage 138

configuring a deduplication storage unit 132

configuring cloud storage 69

D

Deduplication storage unit

Only use the following media servers 134

Use any available media server 134

Disk type 134

Dynamic Host Configuration Protocol (DHCP) 90

E

encryption

properties 119

see also 93

F

Features and functionality 8

FlashBackup policy

Maximum fragment size (storage unit setting) 135

G

glacier vault

back up 42

restore 43

H

host ID-based certificates

deploying with a token 91

deploying without a token 90

host name-based certificates

deploying 89

hotfix 89

J

job ID search in unified logs 161

K

Key Management Service (KMS) 72

L

legacy logging 162

directories 162

locations 162

LIFECYCLE

backup 47

cloud tiering 45

Local cache directory for CloudCatalyst 22, 54, 62

logging

legacy 162

M

Maximum concurrent jobs 134

Maximum fragment size 135

Microsoft Azure

about 50

configuration options 52

configuration options (advanced) 55

requirements 50

mklogdir.bat 162

Monitoring 148

MSDP cloud storage server

properties 118

N

NetBackup

hotfix 89

NetBackup Accelerator

about 136

NetBackup CloudCatalyst

Cloud storage server properties 108

enabling in Cloud Storage Server Configuration

Wizard 22, 54, 62

ESFS_HOST cloud connection property 114

Local cache directory 22, 54, 62

MSDP cloud storage server properties 118

NetBackup CloudStore Service Container. *See*

CloudStore Service Container

NetBackup Key Management Service (NBKMS) 72

NetBackup Scalable Storage 73–74

NetBackup Scalable Storage host properties

unavailable 168

NetBackup Service Layer (NBSL) 89

O

OpenStack Swift

about 57

OpenStack Swift (*continued*)
 configuration options (cloud storage instance) 23, 65
 provider configuration options 60, 63
 proxy settings 65
 requirements 58
 Optimized Synthetic backups
 about 136

P

policies
 changing properties 141
 creating 140
 port number
 CloudStore Service Container 81
 configuring for the CloudStore Service Container 85
 Preferences
 common 113
 encryption 119
 throttling 119
 private clouds
 Amazon S3-compatible cloud providers 30
 properties
 bandwidth 109
 cloud storage server 108
 CloudCatalyst storage server 118
 connection 112
 encryption 119

R

Replication Director
 Policy Configuration Wizard, unsupported 140
 Reporting 148
 requirements 71

S

Scalable Storage host properties 71, 73–74
 Scalable Storage host properties unavailable 168
 Scalable Storage, NetBackup 73–74
 security certificates
 for cloud storage 82
 server
 NetBackupdebug logs 162
 Status Collection Daemon 163
 storage server. *See* cloud storage server
 changing properties for cloud 106

storage unit
 configuring for deduplication 132
 properties for cloud 133
 Storage unit name 133
 Storage unit type 134

T

throttling data transfer rate 72

U

unified logging 158
 format of files 160
 location 158

V

virtual private cloud 32
 vmscd 163
 vmscd directory 163
 VPC 32
 vxlogview command 159
 with job ID option 161

W

wizards
 Policy Configuration 140