

NetBackup™ for OpenStack Administrator's Guide

UNIX, Windows, and Linux

VERITAS™

OpenStack Administrator's Guide

Last updated: 2022-02-23

Legal Notice

Copyright © 2022 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
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

NB.docs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Introduction	6
	Protecting OpenStack data using NetBackup	6
	Backing up OpenStack data	8
	Restoring OpenStack data	9
	NetBackup for OpenStack terminologies	10
Chapter 2	Deploying OpenStack plug-in for NetBackup	12
	About the OpenStack plug-in deployment	12
	Operating system and platform compatibility	13
	License for OpenStack plug-in for NetBackup	13
	Preparing OpenStack	13
	Verifying the deployment of the OpenStack plug-in	14
Chapter 3	Configuring NetBackup for OpenStack	15
	About configuring NetBackup for OpenStack	15
	Managing backup hosts	16
	Whitelisting a NetBackup client on NetBackup master server	20
	Adding OpenStack credentials in NetBackup	20
	Configuring the OpenStack plug-in using the OpenStack configuration file	27
	Configuring NetBackup BigData policy for OpenStack	27
	Creating BigData policy using the NetBackup Administration Console	28
	Using NetBackup Command Line Interface (CLI) to create a BigData policy for OpenStack	29
Chapter 4	Performing backups and restores of OpenStack	32
	About backing up OpenStack data	32
	Backing up OpenStack data	33
	About the metadata information that is captured during the backup	33
	About restoring OpenStack data	34
	Using the Restore Wizard to restore OpenStack data	35

	Using the <code>bprestore</code> command to restore OpenStack data	37
Chapter 5	Troubleshooting	41
	About NetBackup for OpenStack debug logging	41
	Known limitations for OpenStack protection using NetBackup	42
Index		43

Introduction

This chapter includes the following topics:


- [Protecting OpenStack data using NetBackup](#)
- [Backing up OpenStack data](#)
- [Restoring OpenStack data](#)
- [NetBackup for OpenStack terminologies](#)

Protecting OpenStack data using NetBackup

Using the NetBackup Parallel Streaming Framework (PSF), OpenStack data can now be protected using NetBackup.

[Protection support for OpenStack is deployed on CEPH Storage and Backup Host deployment of global admin](#)

The following diagram provides an overview of how OpenStack data is protected by NetBackup.

Figure 1-1 Architectural overview

As illustrated in the above diagram:

- The data is backed up in parallel streams with the help of backup hosts deployed in OpenStack environment. You can select the instances you want to protect within the project. The job processing is accelerated due to multiple backup hosts and parallel streams.
- The communication between the OpenStack and the NetBackup is enabled using the NetBackup plug-in for OpenStack. The plug-in is installed with the NetBackup installation.
- For NetBackup communication, you need to configure a Big Data policy and add the related backup hosts.
- You can configure a NetBackup media server or client as a backup host.
- The backup host can be also used as a finger printing media server for deduplication load balancing. For more information, see [NetBackup Deduplication Guide](#).
- Use the NetBackup media server as a backup host.

For more information:

- See "[Backing up OpenStack data](#)" on page 8.

- See “[Restoring OpenStack data](#)” on page 9.

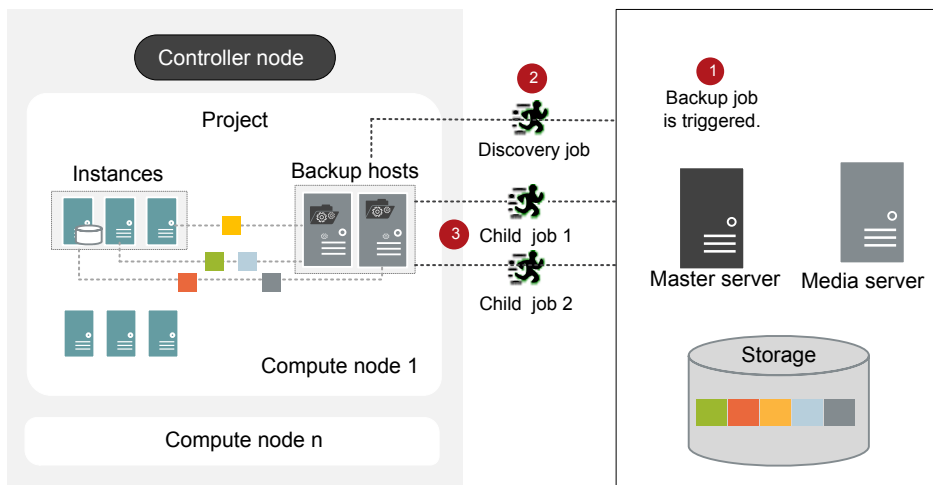
Note: OpenStack deployed in demilitarized zones (DMZ) cannot be protected using this solution.

Backing up OpenStack data

OpenStack data is backed up in parallel streams wherein OpenStack project instances stream data blocks to the NetBackup media server through the backup hosts.

The following diagram provides an overview of the backup flow:

Figure 1-2 Backup flow



As illustrated in the above diagram:

1. A scheduled backup job is triggered from the master server.
2. Backup job for OpenStack data is a compound job. When the backup job is triggered, first a discovery job is run.
3. During discovery, the first backup host performs a discovery to get details of data that needs to be backed up.

- A workload discovery file is created on the backup host. The workload discovery file contains the details of the data that needs to be backed up from the different instances.
 - The backup host uses the workload discovery file and decides how the workload is distributed amongst the backup hosts. Workload distribution files are created for each backup host.
4. Individual child jobs are executed for each backup host. As specified in the workload distribution files, data is backed up.

The compound backup job is not completed until all the child jobs are completed.

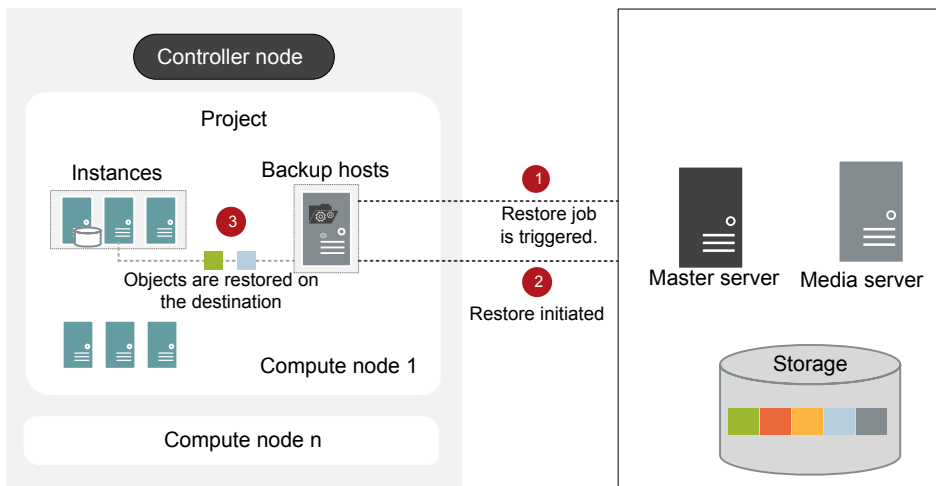
See [“About backing up OpenStack data”](#) on page 32.

Restoring OpenStack data

For restore only one backup host is used.

The following diagram provides an overview of the restore flow.

Figure 1-3 Restore flow



As illustrated in the above diagram:

1. The restore job is triggered from the master server.
2. Backup host is the destination client.

- The objects are restored on the instances of the controller node. New objects are created on the destination.

See [“About restoring OpenStack data”](#) on page 34.

NetBackup for OpenStack terminologies

The following table defines the terms you will come across when using NetBackup for protecting OpenStack.

Table 1-1 NetBackup terminologies

Terminology	Definition
Compound job	<p>A backup job for OpenStack data is a compound job.</p> <ul style="list-style-type: none">The backup job runs a discovery job for getting information of the data to be backed up.Child jobs are created for each backup host that performs the actual data transfer.
Discovery job	<p>When a backup job is executed, first a discovery job is created. The discovery job communicates with the controller node and gathers information of the instances and associated volumes (cinder) that needs to be backed.. At the end of the discovery, the job populates a workload discovery file that NetBackup then uses to distribute the workload amongst the backup hosts.</p>
Child job	<p>For backup, a separate child job is created for each backup host to transfer data to the storage media. A child job can transfer data blocks from multiple OpenStack servers.</p>
Parallel streams	<p>The NetBackup parallel streaming framework allows the instances and associated volumes (cinder) backed up using multiple backup hosts sequentially.</p>
Backup host	<p>The backup host acts as a proxy client. All the backup and restore operations are executed through the backup host.</p> <p>You can configure media servers or clients as a backup host.</p> <p>The backup host is also used as destination client during restores.</p>
BigData policy	<p>The BigData policy is introduced to:</p> <ul style="list-style-type: none">Specify the application type.Allow backing up distributed multi-node environments.Associate backup hosts.Perform workload distribution.

Table 1-1 NetBackup terminologies (*continued*)

Terminology	Definition
Application server	Controller node is referred to as an application server in NetBackup.

Deploying OpenStack plug-in for NetBackup

This chapter includes the following topics:

- [About the OpenStack plug-in deployment](#)
- [Operating system and platform compatibility](#)
- [License for OpenStack plug-in for NetBackup](#)
- [Preparing OpenStack](#)
- [Verifying the deployment of the OpenStack plug-in](#)

About the OpenStack plug-in deployment

The OpenStack plug-in is installed with NetBackup. Review the following topics to complete the deployment.

Table 2-1 OpenStack plug-in deployment

Task	Reference
Pre-requisites and requirements	See “Operating system and platform compatibility” on page 13. See “License for OpenStack plug-in for NetBackup” on page 13.
Preparing the OpenStack	See “Preparing OpenStack ” on page 13.
Verifying the deployment	See “Verifying the deployment of the OpenStack plug-in ” on page 14.
Configuring	See “About configuring NetBackup for OpenStack” on page 15.

Operating system and platform compatibility

With this release, the following are supported:

- Supported OpenStack versions: Mitaka, Newton, Ocata, Pike, and Queens
- Supported authentication types: Simple

For more information, see the [NetBackup Master Compatibility List](#).

License for OpenStack plug-in for NetBackup

Review the following tech note and apply the appropriate license:

https://www.veritas.com/content/support/en_US/article.100040155.html

More information is available on how to add licenses.

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

Preparing OpenStack

Perform the following tasks to prepare OpenStack for NetBackup:

- Update firewall settings so that the backup hosts can communicate with the OpenStack endpoints - Nova, Keystone, Glance, Cinder, and Neutron.
- Ensure that backup hosts can communicate with the NetBackup Master Server.
- The backup host must be deployed in the compute node. See [“Managing backup hosts”](#) on page 16.
- Add the entries of the controller node and the compute node to the `/etc/hosts` file on the associated backup hosts. You must add the hostname in FQDN format or add the appropriate DNS entries in the `/etc/resolv.conf` file.
Or
Add the appropriate DNS entries in the `/etc/resolve.conf` file.
- Use consistent conventions for hostnames of backup hosts, media servers, and master server. For example, if you are using the hostname as **openstack.veritas.com** (FQDN format), use the same everywhere.
- Display name and hostname of the backup hosts should be the same.

Verifying the deployment of the OpenStack plug-in

After you install NetBackup, the

`/usr/opensv/lib/psf-plugins/openstack/libaapipgnopenstack.so` file is deployed.

Configuring NetBackup for OpenStack

This chapter includes the following topics:

- [About configuring NetBackup for OpenStack](#)
- [Managing backup hosts](#)
- [Whitelisting a NetBackup client on NetBackup master server](#)
- [Adding OpenStack credentials in NetBackup](#)
- [Configuring the OpenStack plug-in using the OpenStack configuration file](#)
- [Configuring NetBackup BigData policy for OpenStack](#)

About configuring NetBackup for OpenStack

Table 3-1 Configuring NetBackup for OpenStack

Task	Reference
Adding backup hosts	See “Managing backup hosts” on page 16. If you want to use NetBackup client as a backup host, you need to whitelist the NetBackup client on the master server. See “Whitelisting a NetBackup client on NetBackup master server” on page 20.
Adding OpenStack credentials in NetBackup	See “Adding OpenStack credentials in NetBackup” on page 20.

Table 3-1 Configuring NetBackup for OpenStack (*continued*)

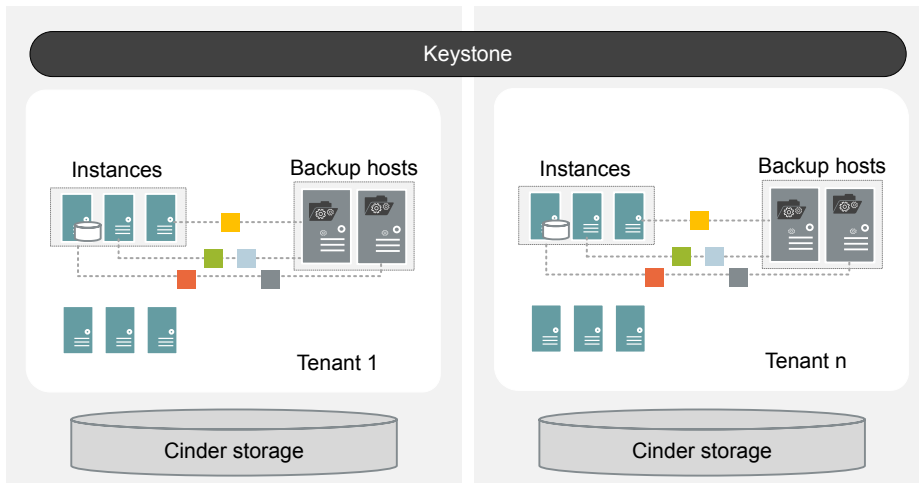
Task	Reference
Configuring the OpenStack plug-in using the OpenStack configuration file	See “ Configuring the OpenStack plug-in using the OpenStack configuration file ” on page 27.
Configuring NetBackup policies for OpenStack plug-in	See “ Configuring NetBackup BigData policy for OpenStack ” on page 27.

Managing backup hosts


The backup host is supported on RHEL and SUSE operating systems. See <http://www.netbackup.com/compatibility>.

The following backup host deployment models are supported for OpenStack protection using NetBackup:

- Local admin
Backup hosts are deployed for each tenant or project that needs protection.



- Global admin
Backup hosts are deployed for a special tenant (backup tenant) that is responsible for performing backup or restore operations for all the other tenants or projects within the OpenStack cluster.



Backup host must not be used as a shared client, wherein the media server or client you are using cannot associate with multiple master servers.

Consider the following before adding a backup host:

- For backup operations, you can add one or more backup hosts. One backup host can manage up to 40 instances, thus add the number of backup hosts accordingly.
- For restore operations, you can add only one backup host.
- Make sure that the backup hosts are communicating with the NetBackup media and master server.

You can add a backup host while configuring BigData policy using either the NetBackup Administration Console or Command Line Interface.

For more information on how to create a policy, see See [“Configuring NetBackup BigData policy for OpenStack”](#) on page 27.

To add a backup host

- 1 In the **Backup Selections** tab, click **New** and add the backup host in the following format:

Backup_Host=<hostname>

For more information on how to create a policy, See [“Configuring NetBackup BigData policy for OpenStack”](#) on page 27.

Alternatively, you can also add a backup host using the following command:

For Windows:

```
bpplinclude PolicyName -add "Backup_Host=hostname"
```

For UNIX:

```
bpplinclude PolicyName -add 'Backup_Host=hostname'
```

For more information, See [“Using NetBackup Command Line Interface \(CLI\) to create a BigData policy for OpenStack ”](#) on page 29.

- 2 As a best practice, add the entries of all the controller node and compute node to the `/etc/hosts` file on all the backup hosts. You must add the host name in FQDN format and add the appropriate DNS entries in the `/etc/resolv.conf` file.

OR

Add the appropriate DNS entries in the `/etc/resolve.conf` file.

- 3 (Optional) If you are using a media server as backup host that is deployed on the OpenStack instance, add that media server to the master server host properties.

To remove a backup host

- 1 In the **Backup Selections** tab, select the backup host that you want to remove.
- 2 Right click the selected backup host and click **Delete**.

Alternatively, you can also remove a backup host using the following command:

For Windows:

```
bpplinclude PolicyName -delete "Backup_Host=hostname"
```

For UNIX:

```
bpplinclude PolicyName -delete 'Backup_Host=hostname'
```

Whitelisting a NetBackup client on NetBackup master server

To use the NetBackup client as a backup host, you must whitelist it. Perform the Whitelisting procedure on the NetBackup master server .

Whitelisting is a security practice used for restricting systems from running software or applications unless these have been approved for safe execution.

Note: Whitelisting is not required for media servers that will be used as backup hosts.

To Whitelist a NetBackup client on NetBackup master server

◆ Run the following command on the NetBackup master server:

■ For UNIX

```
bpsetconfig -h masterserver
bpsetconfig> APP_PROXY_SERVER = clientname.domain.org
bpsetconfig>
UNIX systems: <ctl-D>
```

■ For Windows

```
bpsetconfig -h masterserver
bpsetconfig> APP_PROXY_SERVER = clientname1.domain.org
bpsetconfig> APP_PROXY_SERVER = clientname2.domain.org
bpsetconfig>
Windows systems: <ctl-Z>
```

This command sets the `APP_PROXY_SERVER = clientname` entry in the backup configuration (`bp.conf`) file.

For more information about the `APP_PROXY_SERVER = clientname`, refer to the *Configuration options for NetBackup clients* section in *NetBackup Administrator's Guide, Volume I*

[Veritas NetBackup Documentation](#)

Adding OpenStack credentials in NetBackup

To establish a seamless communication between OpenStack and NetBackup for backup and restore operations, you must add and update OpenStack credentials in the NetBackup master server.

You need to first create a credentials file for storing the Keystone and project information. This file is used as an input when you run the `tpconfig` command to add credentials in NetBackup master server.

You can use the following backup host deployment models to protect OpenStack:

- [Local admin backup host deployment](#)
- [Global admin backup host deployment](#)

For more information, See “[Managing backup hosts](#)” on page 16.

The credential file differs based on the backup host deployment model.

Local admin backup host deployment

In this deployment model, backup hosts are deployed for each tenant or project.

To create a credentials file for storing and entering Keystone and project information

- 1 Login to the NetBackup master server.
- 2 On the OpenStack server, use the following steps to get the information that you need to create the credential file:

- `cat ~/keystonerc_admin`

```
unset OS_SERVICE_TOKEN
export OS_USERNAME=admin1
export OS_PASSWORD='aae1113cd1482a'
export OS_REGION_NAME=RegionOne
export OS_AUTH_URL=http://10.217.34.248:5000/v3
export PS1='[\u@\h \W(keystone_admin)]\$ '
export OS_PROJECT_NAME=admin
export OS_USER_DOMAIN_NAME=Default
export OS_PROJECT_DOMAIN_NAME=Default
export OS_IDENTITY_API_VERSION=3
```

- You required the following variables:

- `OS_USERNAME`
- `OS_PASSWORD`
- `OS_USER_DOMAIN_NAME`
- `OS_AUTH_URL`
- `OS_PROJECT_NAME`
- `OS_PROJECT_DOMAIN_NAME`

- ProjectUUID
 For ProjectUUID: `openstack project list | grep OS_PROJECT_NAME | awk '{print $2}'`
 The output will be ProjectUUID of the PROJECT.
- IPAddress
 Get IP Address of the OpenStack Controller Node. The IP address is used in credential file and in policy as name of client.
- EndPoint
 This value is required for communication. EndPoint examples are internal, public, admin.

■ **Sample credential file format for local admin backup host deployment:**

```
{
  "IPAddress_management_interface": "EndPoint",
  "IPAddress_volume_api_version": "3",
  "IPAddress_ep_keystone": "OS_AUTH_URL",
  "IPAddress_os_access_protocol": "http://",
  "IPAddress_domain_id": "OS_PROJECT_DOMAIN_NAME",
  "IPAddress_auth_sub_url": "auth/tokens",
  "IPAddress_ProjectUUID ":
  ("keystone_username": "OS_USERNAME", "keystone_password": "OS_PASSWORD", "keystone_user_domain_name": "OS_USER_DOMAIN_NAME",
  "project_domain_name": "OS_PROJECT_DOMAIN_NAME",
  "project_name": "OS_PROJECT_NAME", "user_role": "member"},
  "IPAddress_admin":
  ("keystone_username": "OS_USERNAME", "keystone_password": "OS_PASSWORD", "keystone_user_domain_name": "OS_USER_DOMAIN_NAME",
  "project_domain_name": "OS_PROJECT_DOMAIN_NAME",
  "project_name": "OS_PROJECT_NAME", "user_role": "member")
}
```

Sample values for the variables:

```
IPAddress = 10.217.34.248
EndPoint = internal
ProjectUUID = 9c43b3b5d55c414497fb46f7141c604d
OS_AUTH_URL = http://10.217.34.248:5000/v3
OS_PROJECT_DOMAIN_NAME = Default
OS_USERNAME = admin
OS_PASSWORD = aaaa1113cd1482a
OS_USER_DOMAIN_NAME = Default
OS_PROJECT_DOMAIN_NAME = Default
```

Sample credential file using the sample values for local admin backup host deployment:

```
{
  "10.217.34.248_management_interface": "internal",
  "10.217.34.248_volume_api_version": "3",
  "10.217.34.248_ep_keystone": "http://10.217.34.248:5000/v3",
  "10.217.34.248_os_access_protocol": "http://",
  "10.217.34.248_domain_id": "default",
  "10.217.34.248_auth_sub_url": "auth/tokens",
  "10.217.34.248_9c43b3b5d55c414497fb46f7141c604d":
  {"keystone_user": "admin", "keystone_password": "aaaa1113cd1482a", "keystone_user_domain_name": "Default",
   "project_domain_name": "Default", "project_name": "admin"},
  "10.217.34.248_admin":
  {"keystone_user": "admin", "keystone_password": "aaaa1113cd1482a", "keystone_user_domain_name": "Default",
   "project_domain_name": "Default", "project_name": "admin"}
}
```

- Add the credentials file in the `/usr/opensv/var/global` folder on your NetBackup master server.

3 Whitelist the file path of the creds file. Run the following command:

```
bpsetconfig -h masterserver
```

For UNIX: `<ctl-z>`

For Windows: `<ctl-d>`

Note: Whitelisting is not required for media server to be able to use as backup host.

Global admin backup host deployment

In this deployment model, all the backup hosts are part of a single tenant or project.

To create a credentials file for storing and entering Keystone and project information

- 1 Login to the NetBackup master server.
- 2 On the OpenStack server, use the following steps to get the information that you need to create the credential file:
 - You required the following variables:
 - OS_USERNAME

- OS_PASSWORD
- OS_PROJECT_NAME
- OS_PROJECT_DOMAIN_NAME
- ProjectUUID

For ProjectUUID: `openstack project list | grep OS_PROJECT_NAME
 | awk '{print $2}'`

The output will be ProjectUUID of the PROJECT.

- IPAddress

Get IP Address of the OpenStack Controller Node. The IP address is used in credential file and in policy as name of client.

- Sample credential file format for global admin backup host deployment:

```
{
  " IPAddress _g_backup_admin_name": "GA_USERNAME",
  " IPAddress
  _g_backup_admin_domain_name": "GA_PROJECT_DOMAIN_NAME",
  " IPAddress _g_backup_admin_password": "GA_PASSWORD ",
  " IPAddress _g_backup_admin_project_name": "GA_PROJECT_NAME",
  " IPAddress _g_backup_admin_project_id": "ProjectUUID ",
  " IPAddress
  _g_backup_admin_project_domain_name": "GA_PROJECT_DOMAIN_NAME
  ",

  "IPAddress_management_interface": "EndPoint",
  "IPAddress_volume_api_version": "3",
  "IPAddress_ep_keystone": "OS_AUTH_URL",
  "IPAddress_os_access_protocol": "http://",
  "IPAddress_domain_id": "OS_PROJECT_DOMAIN_NAME",
  "IPAddress_auth_sub_url": "auth/tokens",
  "IPAddress_ProjectUUID ":
  ("keystone_user": "OS_USERNAME", "keystone_password": "OS_PASSWORD", "keystone_user_domain_name": "OS_USER_DOMAIN_NAME",
  "project_domain_name": "OS_PROJECT_DOMAIN_NAME",
  "project_name": "OS_PROJECT_NAME", "user_role": "member"),
  "IPAddress_admin":
  ("keystone_user": "OS_USERNAME", "keystone_password": "OS_PASSWORD", "keystone_user_domain_name": "OS_USER_DOMAIN_NAME",
  "project_domain_name": "OS_PROJECT_DOMAIN_NAME",
  "project_name": "OS_PROJECT_NAME", "user_role": "member")
}
```

Sample values for the variables:


```

IPAddress = 10.217.34.248
EndPoint = internal
ProjectUUID = 9c43b3b5d55c414497fb46f7141c604d
OS_AUTH_URL = http://10.217.34.248:5000/v3
OS_PROJECT_DOMAIN_NAME = Default
OS_USERNAME = admin
OS_PASSWORD = aaea1113cd1482a
OS_USER_DOMAIN_NAME = Default
    
```

Sample credential file using the sample values for global admin backup host deployment:

```

{
  "10.217.34.248_g_backup_admin_name": "admin",
  "10.217.34.248_g_backup_admin_domain_name": "Default",
  "10.217.34.248_g_backup_admin_password": "aea1113cd1482a",
  "10.217.34.248_g_backup_admin_project_name": "admin",
  "10.217.34.248_g_backup_admin_project_id": "9a6de296541c4a62891d8ea0b2aeed05",
  "10.217.34.248_g_backup_admin_project_domain_name": "Default",
  "10.217.34.248_management_interface": "internal",
  "10.217.34.248_volume_api_version": "3",
  "10.217.34.248_ep_keystone": "http://10.217.34.248:5000/v3",
  "10.217.34.248_os_access_protocol": "http://",
  "10.217.34.248_domain_id": "default",
  "10.217.34.248_auth_sub_url": "auth/tokens",
  "10.217.34.248_9a6de296541c4a62891d8ea0b2aeed05":
  {"keystone_user": "admin", "keystone_password": "aea1113cd1482a", "keystone_user_domain_name": "Default",
   "project_domain_name": "Default", "project_name": "admin",
   "backuptime_az": "nova"},
  "10.217.34.248_admin":
  {"keystone_user": "admin", "keystone_password": "aea1113cd1482a", "keystone_user_domain_name": "Default",
   "project_domain_name": "Default", "project_name": "admin",
   "backuptime_az": "nova"},
  "10.217.34.248_12c3cbcaf92b4e13a8c3bb4f74efe513":
  {"keystone_user": "demo", "keystone_password": "5a7499ff22f04729", "keystone_user_domain_name": "Default",
   "project_domain_name": "Default", "project_name": "demo",
   "backuptime_az": "nova", "user_role": "member"},
  "10.217.34.248_demo":
  {"keystone_user": "demo", "keystone_password": "5a7499ff22f04729", "keystone_user_domain_name": "Default",
   "project_domain_name": "Default", "project_name": "demo",
   "backuptime_az": "nova", "user_role": "member"}
}
    
```

- Add the credentials file in the `/usr/opensv/var/global` folder on your NetBackup master server.
- 3 Whitelist the file path of the creds file. Run the following command:

```
bpsetconfig -h masterserver
```

For UNIX: `<ctl-z>`

For Windows: `<ctl-d>`

Note: Whitelisting is not required for media server to be able to use as backup host.

To add credentials in NetBackup

- 1 Run `tpconfig` command from the following directory paths:

On UNIX systems, `/usr/opensv/volmgr/bin/`

On Windows systems, `install_path\Volmgr\bin\`

- 2 Run the `./tpconfig -add -application_server_user_id user ID -application_type openstack -application_server IP Address -password password -application_server_conf /path to creds file -requiredport Port Number`

Ensure that the host name of the backup host is the same as the display name of the backup host used in OpenStack.

- 3 Run the `tpconfig -dappservers` command to verify if the NetBackup master server has the OpenStack credentials added.

The following entry is added in the existing global admin entries when the credential file is added.

```
"user_role": "admin"
```

This entry is optional for admin users but required for non-admin users.

You can mix the two backup host deployment models and create a hybrid deployment model. In this hybrid model, you can have a global admin credentials in credential file and few tenants without `member_role` users. In that case, they will be admin of that project.

About the backup admin role

The backup administrator role lets the user run backup and restore jobs. Use this role to create a user who can be the backup administrator of a given tenant or project. You can also use this role to create users of the type global admin.

Note: Backup admin is a recommended role but is not mandatory to protect OpenStack.

Configuring the OpenStack plug-in using the OpenStack configuration file

The backup hosts use the `openstack.conf` file to save the configuration settings of the OpenStack plug-in. You need to manually create the `openstack.conf` file in key-value pair format. You need to create the file on the master server at the `/usr/opensv/var/global/` location. This file is not available by default with the installer.

Note: You must not provide a blank value for any of the parameters, or the backup job fails.

With this release, the following plug-in settings can be configured:

- **Connection retries:** The number of retries per connection during operations between NetBackup and OpenStack. By default the value is 100.
- **Read time out:** The timeout value during read operations in seconds. By default the value is 3 seconds.
- **Snapshot during discovery:** Determines if the snapshot is taken during discovery (true) or during backup. It is recommended to perform snapshots during backup. This parameter is optional and by default is set to false.

Following is an example of the `openstack.conf` file.

```
openstack_connection_retries = in numbers
openstack_read_timeout = in seconds
snapshot_during_discovery = false or true
```

Configuring NetBackup BigData policy for OpenStack

For OpenStack, use the NetBackup **BigData** policy with `openstack` as application type.

You can create **BigData** policy using either the **NetBackup Administration Console** or the **Command Line Interface**.

See [“Creating BigData policy using the NetBackup Administration Console”](#) on page 28.

See [“Using NetBackup Command Line Interface \(CLI\) to create a BigData policy for OpenStack”](#) on page 29.

Creating BigData policy using the NetBackup Administration Console

If you prefer using the **NetBackup Administration Console** for creating BigData policy, you can use either of the following methods:

- Creating a BigData policy using the **Policy Configuration Wizard**
- Creating a BigData policy using the NetBackup **Policies** utility

To create a BigData 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 the type of policy to create:
 - **BigData** policy : A policy to backup **openstack** data
- 4 Select the storage unit type for BigData policy.
- 5 Click **Next** to start the wizard and follow the prompts.
Click **Help** on any wizard panel for assistance while running the wizard.

To create a BigData policy with the NetBackup Policies utility

- 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.
Click **OK**.
- 4 On the **Attributes** tab, select **BigData** as the policy type.
- 5 On the **Attributes** tab, select the storage unit for BigData policy type.

- 6 On the **Schedules** tab, click **New** to create a new schedule.
 You can create a schedule for a **Full Backup**, for your BigData policy. Once you set the schedule, OpenStack data is backed up automatically as per the set schedule without any further user intervention.
- 7 On the **Clients** tab, enter the IP address or the host name of the NameNode.
- 8 On the **Backup Selections** tab, enter the following parameters and their values as shown:
 - *Application_Type=openstack*
 The parameter values are case-sensitive.
 - *Backup_Host=hostname*
 The backup host must be a Linux computer. The backup host can be a NetBackup client or a media server.
 You can specify multiple backup hosts.
 - Instance to back up.
 You can specify multiple file paths.

Note: Instance name used for backup selection while defining BigData Policy with *Application_Type=openstack* must not contain space or special character in their names.

- 9 Click **OK** to save the changes.

Using NetBackup Command Line Interface (CLI) to create a BigData policy for OpenStack

You can also use the CLI method to create a BigData policy for OpenStack.

To create a BigData policy using NetBackup CLI method

- 1 Log on as an Administrator.
- 2 Navigate to `/usr/opensv/netbackup/bin/admincmd`.
- 3 Create a new BigData policy using the default settings.

```
bpolicynew policyname
```
- 4 View the details about the new policy using the `-L` option.

```
bpplinfo policyname -L
```
- 5 Modify and update the policy type as **BigData**.

```
bpplinfo PolicyName -modify -v -M MasterServerName -pt BigData
```

6 Specify the *Application_Type* as **openstack**.

For Windows:

```
bpplinclude PolicyName -add "Application_Type=openstack"
```

For UNIX:

```
bpplinclude PolicyName -add 'Application_Type=openstack'
```

Note: The parameter values for *Application_Type=openstack* are case-sensitive.

7 Specify the backup host on which you want the backup operations to be performed for OpenStack.

For Windows:

```
bpplinclude PolicyName -add "Backup_Host=hostname"
```

For UNIX:

```
bpplinclude PolicyName -add 'Backup_Host=hostname'
```

Note: The backup host must be a Linux computer. The backup host can be a NetBackup client or a media server or a master server.

8 Specify the OpenStack directory or folder name that you want to backup.

For Windows:

```
bpplinclude PolicyName -add '/combination of project and instance'
```

For UNIX:

```
bpplinclude PolicyName -add '/combination of project and instance'
```

Supported combinations:

- /Project_name/Instance_name
- /Project_name/Instance_ID
- /Project_ID/Instance_name
- /Project_ID/Instance_ID

Note: Instance name used for backup selection while defining BigData Policy with *Application_Type=openstack* must not contain space or special character in their names.

- 9** Modify and update the policy storage type for BigData policy.

```
bpplinfo PolicyName -residence STUName -modify
```

- 10** Specify the IP address or the host name of the controller node for adding the client details.

For Windows:

```
bpplclients PolicyName -M "MasterServerName" -add  
"OpenStackServerHMaster" "Linux" "RedHat"
```

For UNIX:

```
bpplclients PolicyName -M 'MasterServerName' -add  
'OpenStackServerHMaster' 'Linux' 'RedHat'
```

- 11** Assign a schedule for the created BigData policy as per your requirements.

```
bpplsched PolicyName -add Schedule_Name -cal 0 -rl 0 -st  
sched_type -window 0 0
```

Here, *sched_type* value can be specified as follows:

For *sched_typeonly* **FULL** is supported.

Once you set the schedule, OpenStack data is backed up automatically as per the set schedule without any further user intervention.

- 12** Alternatively, you can also perform a manual backup for OpenStack data.

For performing a manual backup operation, execute all the steps from Step 1 to Step 11.

- 13** For a manual backup operation, navigate to `/usr/opensv/netbackup/bin`

Initiate a manual backup operation for an existing BigData policy using the following command:

```
bpbackup -i -p PolicyName -s Schedule_Name -S MasterServerName  
-t 44
```

Here, `-p` refers to policy, `-s` refers to schedule, `-S` refers to master server, and `-t 44` refers to BigData policy type.

Performing backups and restores of OpenStack

This chapter includes the following topics:

- [About backing up OpenStack data](#)
- [About restoring OpenStack data](#)

About backing up OpenStack data

Use the **NetBackup, Backup, Archive, and Restore** console to manage backup operations.

Table 4-1 Backing up OpenStack data

Task	Reference
Process understanding	See “Backing up OpenStack data” on page 8.
Backing up OpenStack	See “Backing up OpenStack data” on page 33.
Troubleshooting tips	For discovery and cleanup related logs, review the following log file on the first backup host that triggered the discovery. <code>/usr/openv/netbackup/logs/nbaapidiscv</code> For data transfer related logs, search for corresponding backup host (using the hostname) in the log files on the master server. See “About NetBackup for OpenStack debug logging” on page 41.

Backing up OpenStack data

You can either schedule a backup job or run a backup job manually. See, [NetBackup Administrator's Guide, Volume I](#)

For overview of the backup process, See "[Backing up OpenStack data](#)" on page 8.

The backup process comprises of the following stages:

1. Pre-processing: In the pre-processing stage, the first backup host that you have configured with the BigData policy, triggers the discovery.
2. Data transfer: During the data transfer process, one child job is created for each backup host.
3. Post-processing: As a part of the post-processing, NetBackup cleans up the snapshots from the OpenStack environment.

Note: To avoid double licensing charge, in the backup selection if there is combination of human readable name (HRN) and UUID for the instance from the same master server, provide either of one in the backup selection.

About the metadata information that is captured during the backup

NetBackup captures the following metadata information for OpenStack instance and volumes:

- VolumeType:
This field tells about backend used by cinder for a particular volume. For example, CEPH, iSCSI, etc.
- DeleteOnTermination:
If set to true, volume is deleted when instance is deleted. If set to false, volume is not deleted when instance is deleted. This information is preserved during restore using OpenStack plugin.
 - This field is valid and captured only if instance we are protecting is booting from volume.
 - This field is not captured if instance we are protecting is booting from image. NetBackup preserves "volume type" and "delete_on_termination" for attached volumes irrespective of instance type (BootFromImage or BootFromVolume).
- Apart from above two we are also capturing "Key Name" and "Properties" fields in instance metadata during backup process and are preserved during restore.
- Keypair:
NetBackup preserves and restores keypair only if the keypair is available on OpenStack project at the time of restore otherwise it will restore instance with

default key_name i.e. None. NetBackup does not create keypair for instance recovery.

- **Meta Property:**
 We have put restriction on property field length. The number of characters should be less than 255 chars.

About restoring OpenStack data

Use the **NetBackup, Backup, Archive, and Restore** console to manage restore operations.

To restore OpenStack data, consider following:

- Use the Backup, Archive, and Restore console to initiate OpenStack data restore operations. This interface lets you select the NetBackup server from which the objects are restored and the client whose backup images you want to browse. Based upon these selections, you can browse the backup image history, select individual items and initiate a restore.
- The restore browser is used to display OpenStack objects. A hierarchical display is provided where objects can be selected for restore. The objects (OpenStack instances and attached volumes) are displayed by expanding an individual directory.
- An administrator can browse for and restore OpenStack instances with attached volumes.

Table 4-2 Restoring OpenStack data

Task	Reference
Process understanding	See “Restoring OpenStack data” on page 9.
Restoring OpenStack data to the original location or to an alternate location	<ul style="list-style-type: none"> ■ See “Using the Restore Wizard to restore OpenStack data” on page 35. ■ See “Using the <code>bprestore</code> command to restore OpenStack data” on page 37.
Troubleshooting tips	See “About NetBackup for OpenStack debug logging” on page 41.

Using the Restore Wizard to restore OpenStack data

This topic describes how to use the Restore Wizard to restore OpenStack data on the same OpenStack cluster.

To use the Restore Wizard to perform a restore

- 1 Open the **Backup, Archive, and Restore** interface.
- 2 Select the appropriate date range to restore the complete data set.
- 3 In the **Browse** directory, specify the root directory (“/”) as the path to browse.
- 4 From the File menu (Windows) or Actions menu (UNIX), choose **Specify NetBackup Machines and Policy Type**.
- 5 On the **Specify NetBackup Machines and Policy Type** wizard, enter the source and destination details for restore.
 - Specify the OpenStack controller node as the source for which you want to perform the restore operation.
From the **Source client for restores** list, select the required controller node.
 - Specify the backup host as the destination client.
From the **Destination client for restores** list, select the required backup host.
 - On the **Specify NetBackup Machines and Policy Type** wizard, enter the policy type details for restore.
From the **Policy type for restores** list, choose **BigData** as the policy type for restore.
Click **Ok**.
- 6 Go to the **Backup History** and select the backup images that you want to restore.
- 7 In the **Directory Structure** pane, expand the **Directory**.
All the data and metadata files and folders under the directory are displayed in the **Contents of Selected Directory** pane.
- 8 (Optional) You can modify the metadata related to instance and volume.
- 9 In the **Contents of Selected Directory** pane, select the check box for the OpenStack files that you want to restore.
- 10 Click **Restore**.
- 11 In the **Restore Marked Files** dialog box, select the destination for restore as per your requirement.

- Select **Restore everything to its original location** if you want to restore your files to the same location where you performed your backup.
- Select **Restore everything to a different location** if you want to restore your files to a location which is not the same as your backup location.

12 Click **Start Restore**.

13 Verify the restored instances or volumes.

Note: If the default instance name is used, the instance is restored with `<SourceInstanceName>_RESTORED name`.

14 A new object is created on destination location.

Modifying the metadata related to instance or volume before restore

During restore you can modify the following metadata related to an instance:

- Auto recover
- Flavor
- Instance name
- Restore availability zone
- Size
- State

During restore you can modify the following metadata related to volume:

- Volume size
- Volume name
- Volume availability zone

To modify the metadata before restore

1 In the **Directory Structure** pane, expand the Directory.

All the subsequent data and meta files and folders under the directory are displayed in the **Contents of Selected Directory** pane.

2 Select the instance that you want to restore.

3 Click the selected Metadata directory, and in the **Contents of Selected Directory** pane, reselect (deselect and then select) the metadata that you want to modify.

- 4 Click **Restore**.
- 5 In the **Restore Marked Files** dialog box, select Restore individual directories and files to different locations.
- 6 For every metadata value that you want to change, select the value, click **Change Selected Destination(s)**, and in the **Destination** field modify the metadata value at the end of the URL.

Using the `bprestore` command to restore OpenStack data

To restore OpenStack data on the same location as your backup location

- 1 Log on as an Administrator or root user based on windows or UNIX system respectively.
- 2 Run the following command on the NetBackup master server by providing appropriate values:

```
bprestore -S master_server -D backup_host -C client -t 44 -L  
progress_log -f listfile
```

Where,

```
-S master_server
```

Specifies the name of the NetBackup master server.

```
-D backup host
```

Specifies the name of the backup host.

```
-C client
```

Specifies a controller node as a source to use for finding backups or archives from which to restore files. This name must be as it appears in the NetBackup catalog.

```
-f listfile
```

Specifies a file (`listfile`) that contains a list of files to be restored and can be used instead of the file names option. In `listfile`, each file path must be on a separate line.

```
-L progress_log
```

Specifies the name of whitelisted file path in which to write progress information.

```
-t 44
```

Specifies BigData as the policy type.

To perform redirected restore for OpenStack

- 1 Modify the values for *rename_file* and *listfile* as follows:

Parameter	Value
<i>rename_file</i>	<p>ALT_APPLICATION_SERVER=<alternate name node></p> <p>The rename file must also contain the changed <code>NetworkID</code> entry.</p> <p>For example, change:</p> <pre style="margin-left: 20px;">/project_name/instance_name/Metadata/NetworkID=value</pre> <p>to</p> <pre style="margin-left: 20px;">/destination_project_name/instance_name/Metadata/NetworkID=value</pre> <p>Note: Alternate restore is supported only to same project name.</p>
<i>listfile</i>	<p>List of all the OpenStack files and metadata files to be restored.</p> <p>The file paths must start with / (slash).</p>

- 2 To fetch the credentials information for the alternate OpenStack controller:
 - Add a `tpconfig` entry for the new OpenStack controller.
 - Change the name of the generated encrypted file in `/usr/opensv/var/global` to match the name of the source client encrypted credentials file.

For example, if `hostname1.conf` is the encrypted source client file and `hostname2.conf` is the encrypted alternate client file. You need to rename

`hostname2.conf` to `hostname1.conf` before running the `bprestore` command.

- 3** Run the `bprestore -S master_server -D backup_host -C client -R rename_file -t 44 -L progress_log -f listfile` command on the NetBackup master server using the modified values for the mentioned parameters in step 1.

Where,

`-S master_server`

Specifies the name of the NetBackup master server.

`-D backup_host`

Specifies the name of the backup host.

`-C client`

Specifies an OpenStack controller as a source to use for finding backups or archives from which to restore files. This name must be as it appears in the NetBackup catalog.

`-f listfile`

Specifies a file (listfile) that contains a list of files to be restored and can be used instead of the file names option. In listfile, each file path must be on a separate line.

`-L progress_log`

Specifies the name of whitelisted file path in which to write progress information.

`-t 44`

Specifies BigData as the policy type.

`-R rename_file`

Specifies the name of a file with name changes for alternate-path restores.

Use the following form for entries in the rename file:

```
ALT_APPLICATION_SERVER=<Application Server Name>
```

To change the volume type at the destination OpenStack environment, add the following line with the proper values for old and new volume types:

```
change /project/instance/Metadata/Src_VolumeType to  
/project/instance/Metadata/Dest_VolumeType
```

Note: Ensure that you have whitelisted all the file paths such as `<rename_file_path>`, `<progress_log_path>` that are already not included as a part of NetBackup install path.

Troubleshooting

This chapter includes the following topics:

- [About NetBackup for OpenStack debug logging](#)
- [Known limitations for OpenStack protection using NetBackup](#)

About NetBackup for OpenStack debug logging

NetBackup maintains process-specific logs for the various processes that are involved in the backup and restore operations. Examining these logs can help you to find the root cause of an issue.

These log folders must already exist for logging to occur. If these folders do not exist, you must create them.

The log folders reside on the following directories

- On Windows: `install_path\NetBackup\logs`
- On UNIX or Linux: `/usr/opensv/netbackup/logs`

Table 5-1 NetBackup logs related to OpenStack

Log Folder	Messages related to	Logs reside on
<code>install_path/NetBackup/logs/bpVMutil</code>	Policy configuration	Master server
<code>install_path/NetBackup/logs/nbaapidisv</code>	BigData framework, discovery, and OpenStack configuration file logs	Backup host

Table 5-1 NetBackup logs related to OpenStack *(continued)*

Log Folder	Messages related to	Logs reside on
install_path/NetBackup/logs/bpbarm	Policy validation, backup, and restore operations	Media server
install_path/NetBackup/logs/bpbkar	Backup	Backup host
install_path/NetBackup/logs/tar	Restore and OpenStack configuration file	Backup host

For more details, refer to the [NetBackup Logging Reference Guide](#).

Known limitations for OpenStack protection using NetBackup

The following table lists the known limitations for OpenStack protection using NetBackup:

Table 5-2 Known limitations

Limitation	Workaround
NetBackup does not clean up the OpenStack volumes if the backup operation fails while downloading the volume. The OpenStack dashboard displays that the volume has an error state.	For OpenStack, NetBackup deletes a volume only if the volume has a state other than error. Workaround: To delete a volume that has the error state, contact the OpenStack Administrator to delete the volumes manually.
OpenStack backup job fails with error 6619 in the following scenarios: <ul style="list-style-type: none"> ■ The OpenStack services go down during the backup. ■ The OpenStack is not responding or is slow in response. ■ Connection is reset during the backup. 	Workaround: Restart the OpenStack service and reinitiate the backup.

Index

A

Adding
 backup host 16

B

Backup 32–33
backup 8
BigData policy
 Command Line Interface 29
 NetBackup Administration Console 28

C

compatibility
 supported operating system 13
configuration file 27
configuring 15

D

deployment 12

I

Installing
 verifying 14

L

License 13

N

NetBackup
 debug logging 41

O

OpenStack credentials
 adding 20
overview
 backup 6
 configuration 6
 deployment 6

overview (*continued*)

 installation 6
 restore 6

P

parallel streaming framework 6
policies
 configuring 27
Preparing 13

R

Removing
 backup host 16
Restore 34
 bprestore command 37
restore 9

T

terms 10

W

Whitelisting
 backuphost 20