

NetBackup™ for OpenStack Administrator's Guide

Release 11.0

NetBackup™ for OpenStack Administrator's Guide

Last updated: 2025-03-04

Legal Notice

Copyright © 2025 Cohesity, Inc. All rights reserved.

Cohesity, Veritas, the Cohesity Logo, Veritas Logo, Veritas Alta, Cohesity Alta, and NetBackup for OpenStack are trademarks or registered trademarks of Cohesity, Inc. 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 Cohesity 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 Cohesity 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 Cohesity, Inc. 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. Cohesity, Inc. 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 Cohesity 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.

Cohesity, Inc.
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 Cohesity 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 Cohesity 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 Cohesity 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	9
	About NetBackup for OpenStack	9
	NetBackup for OpenStack Architecture	10
	Backup as a Service	10
	Main Components	11
	Service Endpoints	12
	Network topology	13
	NetBackup for OpenStack ports	14
Chapter 2	Deploying NetBackup for OpenStack	15
	Requirements	15
	System requirements for NetBackup for OpenStack virtual machine	16
	NetBackup for OpenStack network considerations	17
	Existing endpoints in OpenStack	17
	OpenStack endpoints required by NetBackup for OpenStack	17
	Recommendation: Provide access to all OpenStack Endpoint types	18
	Backup target access required by NetBackup for OpenStack	18
	Example of a typical NetBackup for OpenStack network integration	19
	Other examples of NetBackup for OpenStack network integrations	20
	Preparing the installation	22
	Tenant quotas	22
	NetBackup for OpenStack Cluster	23
	Spinning up the NetBackup for OpenStack virtual machine	23
	Creating the cloud-init image	23
	Spinning up the NetBackup for OpenStack appliance	25
	Uninstalling cloud-init after first start	25
	About NetBackup for OpenStack backup target types	26
	Installing NetBackup for OpenStack Components	26
	Installing on RHOSP	27

Installing on Ansible OpenStack Ussuri	35
Installing on Kolla	42
Configuring NetBackup for OpenStack	55
Details needed for the NetBackup for OpenStack Appliance	56
Advanced settings	59
Starting the configurator	61
Resource throttling in NetBackup for OpenStack	62
Post Installation Health-Check	63
Verify the NetBackup for OpenStack Appliance services are up	63
Check the NetBackup for OpenStack pacemaker and NGINX cluster	65
Verify API connectivity of the NetBackup for OpenStack Appliance	66
Verify that the nbosdm services are up and running	66
Uninstalling NetBackup for OpenStack	67
Uninstalling from RHOSP	67
Uninstalling from Ansible OpenStack	73
Uninstalling from Kolla Openstack	78
Install nbosjm CLI client	81
About log rotation in NetBackup for OpenStack	82
Upgrading NetBackup for OpenStack	86
Deleting the orphaned snapshots	88

Chapter 3 Configuring NetBackup OpenStack Appliance

Reconfigure the NetBackup for OpenStack cluster	89
Configuring the NetBackup primary server details	90
Change NetBackup for OpenStack dashboard password	90
Reset NetBackup for OpenStack dashboard password	91
Downloading the NetBackup for OpenStack logs	91
Updating the API key	91
Uploading the API certificate	92

Chapter 4 Configuring NetBackup primary server

License for OpenStack plug-in for NetBackup	93
About launching the OpenStack Horizon UI from the NetBackup web UI	93
Adding the OpenStack Horizon instance on NetBackup web UI	94
Creating the custom role for NetBackup for OpenStack administrator	94

	Launching the Horizon UI from the NetBackup web UI	95
	Configuring the NBOSVM service principal	95
	About NetBackup for OpenStack protection plan	99
	About auto image replication in NetBackup for OpenStack	99
	Configuring the AIR in NetBackup for OpenStack	100
Chapter 5	NetBackup for OpenStack protections	103
	About protections	103
	List of protections	103
	Create a protection	104
	Protection overview	106
	Edit a protection	107
	Delete a protection	108
	Unlock a protection	109
Chapter 6	Performing snapshots, backups, and restores of OpenStack	110
	About recovery points	111
	List of recovery points	111
	Creating a snapshot	112
	Snapshot and backup overview	113
	Expire recovery points	115
	Cleaning up the volume snapshots	115
	About restores	116
	About restoring the multi-attach volumes	116
	List of Restores	116
	Restores overview	117
	Delete a restore	119
	Cancel a restore	120
	One-click restore	120
	Selective restore	121
	In-place restore	122
	Required restore.json file for CLI	123
	General required information	125
	Selective restore required information	126
	In-place restore required information	130
	About backup mount	132
	Creating a file recovery manager instance	132
	Mounting a backup copy	134
	Accessing the file recovery manager	135
	Identifying mounted backups	135
	Unmounting a backup	136

	About schedules	137
	Enabling or disabling a schedule	137
	Modifying a schedule	137
	About activating the email notifications	138
Chapter 7	Performing Backup Administration tasks	139
	NBOS Backup Admin Area	139
	Access the NBOS Backup Admin area	139
	Configuring the email settings	142
	Enabling or disabling a job scheduler	144
	Protection plan	145
	List the available protection plans	145
	Subscribe a project to a protection plan	146
	Managing the trusts	146
	Policy import and migration	147
	Importing the policies	148
	Orphaned policies	149
Chapter 8	Disaster recovery	151
	About disaster recovery in NetBackup for OpenStack	151
Chapter 9	Troubleshooting	154
	General Troubleshooting Tips	155
	What is happening where	155
	Everything on the Backup Target happens as the user nova	156
	NetBackup for OpenStack Trustee Role	157
	OpenStack Quotas	157
	Ephemeral disk backup	157
	Using the nbosjm CLI tool on the NetBackup for OpenStack Appliance	157
	Health check of NetBackup for OpenStack	158
	On the NetBackup for OpenStack Cluster	158
	The nbosdmapi service	162
	The nbosdm service	163
	Important log files	163
	On the NetBackup for OpenStack Nodes	163
	NetBackup for OpenStack data mover service logs on RHOSP	164
	NetBackup for OpenStack data mover service logs on Ansible	165
	OpenStack	165
	NetBackup for OpenStack data mover service logs on Kolla	166

Troubleshooting NBOSDM container in offline state due to unavailable mount point	166
After restore of the Windows instance, the disk is in an offline state	167
Selective restore from snapshot copy fails	168
A backup fails due to an old nova ID in the universal share path	168
Using the NetBackup support utility in NetBackup for OpenStack	169
Cannot create volumes if the metadata size for physical volume and volume group is small	170
NBOSVM configuration fails if DNS server cannot resolve IP address or IP address is wrong	170
Error when storage unit is created with multiple storage servers	170
Snapshot job fails if the OpenStack image is not accessible to the OpenStack user	171
One-click restore fails if the subnet attached to the instance is not accessible to the OpenStack user	171
The NBOSVM configurator UI does not detect the primary server	172
A recovery point name is updated to a default name	172
NBOS Backups and NBOS Backup Admin tabs disappear from Horizon UI after stack is updated	172
The protection creation fails on the Horizon UI	173
The NetBackup for OpenStack services do not start after NBOSVM is restarted	173
The NBOSVM is not able to communicate with the nbosdmapi on the controller node	173
Troubleshooting the OpenStack Keystone authentication failure	174
Index	175

Introduction

This chapter includes the following topics:

- [About NetBackup for OpenStack](#)
- [NetBackup for OpenStack Architecture](#)

About NetBackup for OpenStack

NetBackup for OpenStack is a native OpenStack service that provides policy-based comprehensive backup and recovery for OpenStack workloads. The solution captures point-in-time workloads (Application, OS, Compute, Network, Configurations, Data, and Metadata of an environment) as full or incremental backups. These backups are held in NetBackup universal share with MSDP, and can be duplicated to NetBackup supported target storages. With NetBackup for OpenStack and its single click recovery, organizations can improve recovery time objectives (RTO) and recovery point objectives (RPO). With NetBackup for OpenStack, IT departments are enabled to fully deploy OpenStack solutions and provide business assurance through enhanced data retention, protection, and integrity.

With the use of NetBackup for OpenStack's VAST (Virtual Snapshot Technology), enterprise IT and the cloud service providers can now deploy backup and disaster recovery as a service to prevent data loss or data corruption through point-in-time snapshots and seamless one-click recovery. NetBackup for OpenStack takes point-in-time backup of the entire workload consisting of compute resources, network configurations, and storage data as one unit. It also takes the incremental backups that only capture the changes that were made since the last backup. Incremental backups save time and storage space as the backup only includes changes since the last backup. The benefits of using VAST for backup and restore can be summarized as follows:

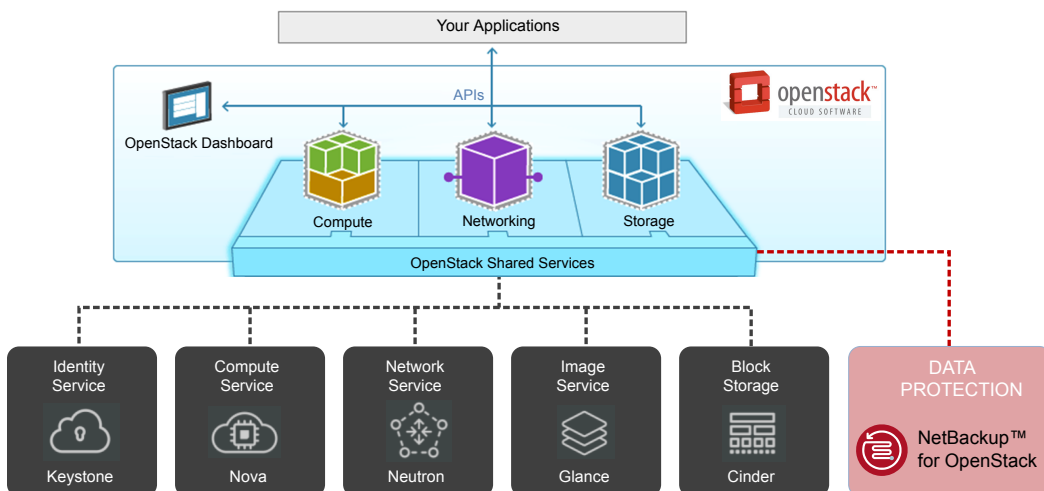
- Efficient capture and storage of snapshots. Since our full backups only include the data that is committed to storage volume and the incremental backups only include changed blocks of data since the last backup, our backup processes are efficient and stores backup images efficiently on the backup media.
- Faster and reliable recovery. When your applications become complex that snap multiple virtual machines and storage volumes, our efficient recovery process brings your application from zero to operational with the click of a button.
- Through policy and automation lower the total cost of ownership. Our tenant-driven backup process and automation eliminates the need for dedicated backup administrators, and improves your total cost of ownership.

NetBackup for OpenStack Architecture

Backup as a Service	See “Backup as a Service” on page 10.
Main components	See “Main Components” on page 11.
Service endpoints	See “Service Endpoints” on page 12.
Network topology	See “Network topology” on page 13.

Backup as a Service

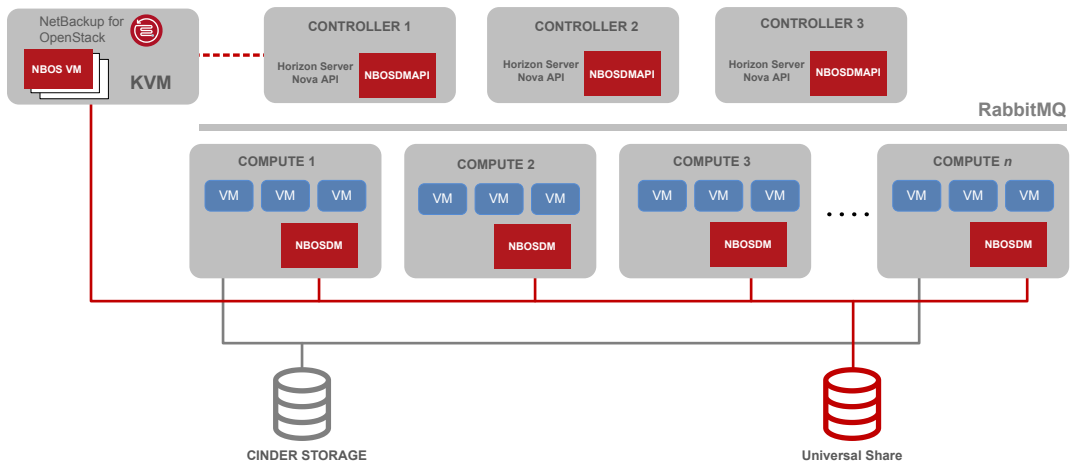
Figure 1-1 Data protection project providing Backup as a Service



NetBackup for OpenStack is an add-on service to OpenStack cloud infrastructure and provides backup and disaster recovery functions for tenant policies. NetBackup for OpenStack is very similar to other OpenStack services including Nova, Cinder, Glance, and adheres to all tenets of OpenStack. This service is a stateless service that scales with your cloud.

Main Components

Figure 1-2 NetBackup for OpenStack architecture overview

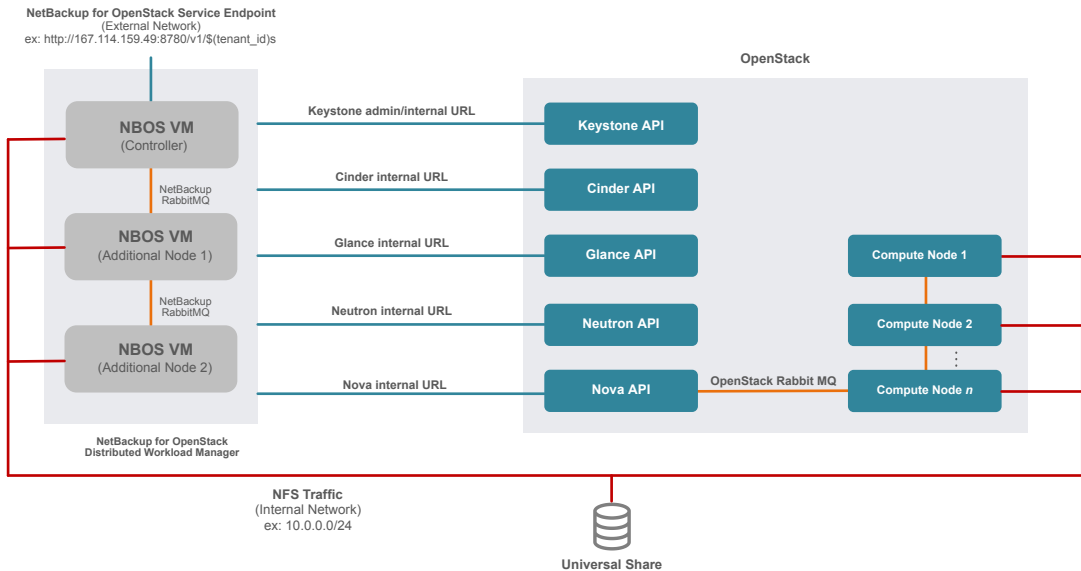


NetBackup for OpenStack has four main software components:

1. NetBackup for OpenStack ships as a QCOW2 image. User can instantiate one or more virtual machines from the QCOW2 image on standalone KVM boxes.
2. NetBackup for OpenStack datamover API (NBOSDMPAPI) is a python module that is installed on all OpenStack controller nodes where the nova-api service is running.
3. NetBackup for OpenStack datamover (NBOSDM) is a python module that is installed on every OpenStack compute nodes
4. NetBackup for OpenStack horizon plug-in is installed as an add-on to horizon servers. This module is installed on every server that runs horizon service.

Service Endpoints

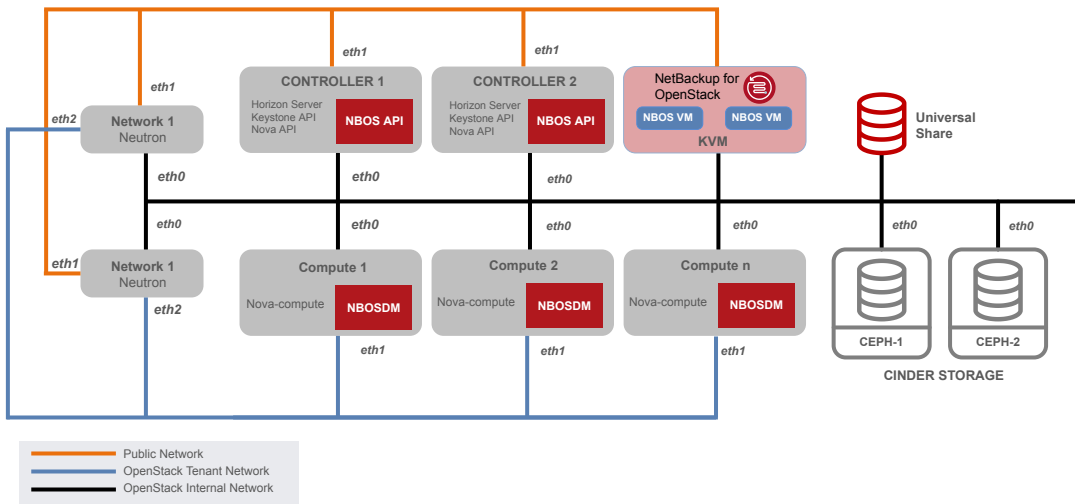
Figure 1-3 Service endpoints overview



NetBackup for OpenStack is both a provider and consumer into the OpenStack ecosystem. It uses other OpenStack services such as nova, cinder, glance, neutron, and keystone and provides its own service to OpenStack tenants. To accommodate all possible OpenStack deployments, NetBackup for OpenStack can be configured to use either public URLs or internal URLs of services. Likewise NetBackup for OpenStack provides its own public, internal, and admin URLs.

Network topology

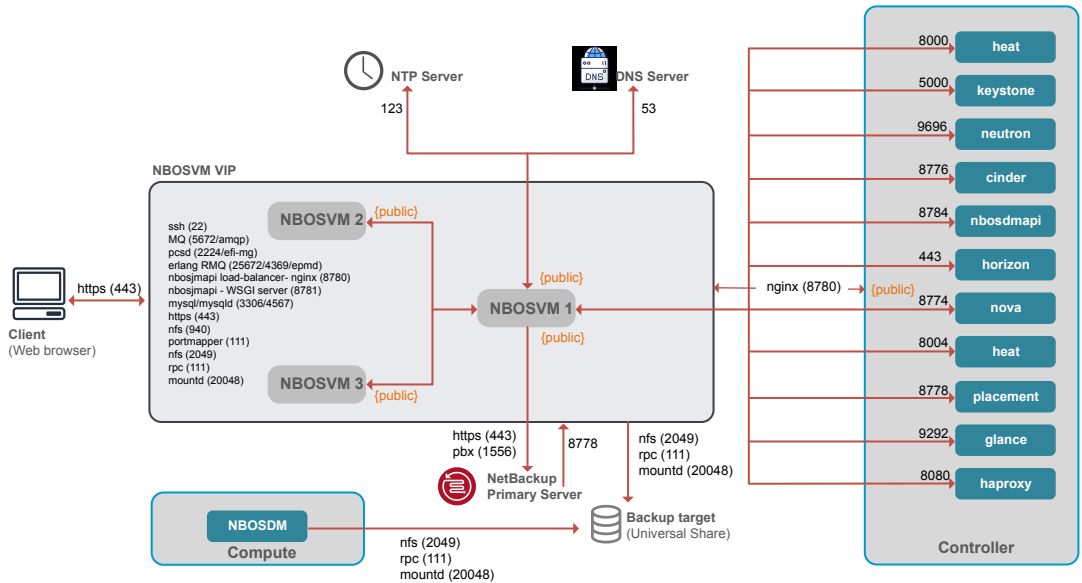
Figure 1-4 Example network topology



This figure represents a typical network topology. NetBackup for OpenStack exposes its public URL endpoint on the public network and NetBackup for OpenStack virtual appliances. The datamovers typically use either the internal network or a dedicated backup network for storing and retrieving backup images from the backup store.

NetBackup for OpenStack ports

Figure 1-5 NetBackup for OpenStack ports



Deploying NetBackup for OpenStack

This chapter includes the following topics:

- [Requirements](#)
- [NetBackup for OpenStack network considerations](#)
- [Preparing the installation](#)
- [Spinning up the NetBackup for OpenStack virtual machine](#)
- [About NetBackup for OpenStack backup target types](#)
- [Installing NetBackup for OpenStack Components](#)
- [Configuring NetBackup for OpenStack](#)
- [Resource throttling in NetBackup for OpenStack](#)
- [Post Installation Health-Check](#)
- [Uninstalling NetBackup for OpenStack](#)
- [Install nbosjm CLI client](#)
- [About log rotation in NetBackup for OpenStack](#)
- [Upgrading NetBackup for OpenStack](#)

Requirements

NetBackup for OpenStack has four main software components:

1. NetBackup for OpenStack ships as a QCOW2 image. You can instantiate one or more virtual machines from the QCOW2 image on standalone KVM boxes.
2. NetBackup for OpenStack Datamover API is a python module that is an extension to Nova API service. This module is installed on all OpenStack controller nodes.
3. NetBackup for OpenStack datamover is a python module that is installed on every OpenStack compute node.
4. NetBackup for OpenStack horizon plug-in is installed as an add-on to horizon servers. This module is installed on every server that runs horizon service.

See [“System requirements for NetBackup for OpenStack virtual machine”](#) on page 16.

See [“Software Requirements ”](#) on page 16.

System requirements for NetBackup for OpenStack virtual machine

The NetBackup for OpenStack virtual machine is delivered as a QCOW2 image, which gets attached to a virtual machine.

Cohesity supports only KVM-based hypervisors.

Note: The NetBackup for OpenStack virtual machine is not supported as an instance inside NetBackup for OpenStack.

The recommended size of the virtual machine for the NetBackup for OpenStack appliance is:

Resource	Value
----------	-------

vCPU	8
RAM	24 GB

The QCOW2 image itself defines the 40GB disk size of the virtual machine.

If the NetBackup for OpenStack virtual machine database or log files get larger than 40GB disk, contact or open a ticket with Cohesity customer support to attach another drive to the NetBackup for OpenStack virtual machine.

Software Requirements

NetBackup for OpenStack is tested and verified.

Software	Version
Red Hat Enterprise Linux	8.9
Virsh	libvirt 2.0.0 and later
QEMU	2.0.0 and later
QEMU disk image utility (qemu-img)	2.6.0 and later

NetBackup for OpenStack network considerations

NetBackup for OpenStack integrates natively with OpenStack. NetBackup for OpenStack communicates completely through APIs using the OpenStack endpoints. NetBackup for OpenStack also generates its own OpenStack endpoints. In addition, is the NetBackup for OpenStack appliance and the compute nodes writing to and reading from the backup target. These points affect the network planning for the NetBackup for OpenStack installation.

Existing endpoints in OpenStack

OpenStack knows three types of endpoints:

- Public endpoints
- Internal endpoints
- Admin endpoints

Each of these endpoint types is designed for a specific purpose. Public endpoints are used by the OpenStack users to work with OpenStack. Internal endpoints are used by the OpenStack services to communicate with each other. Admin endpoints are used by OpenStack administrators.

Out of these three endpoint types, only the admin endpoint sometimes contains APIs, which are not available on any other endpoint type.

To learn more about OpenStack endpoints, visit the official OpenStack documentation.

OpenStack endpoints required by NetBackup for OpenStack

NetBackup for OpenStack communicates with all services of OpenStack on a defined endpoint type. The endpoint type NetBackup for OpenStack uses to communicate with OpenStack is decided during the configuration of the NetBackup for OpenStack appliance. NetBackup for OpenStack supports public endpoints and internal endpoints.

An exception: The NetBackup for OpenStack appliance always requires access to the keystone admin endpoint.

The following network requirements can be identified this way:

- NetBackup for OpenStack appliance needs access to the keystone admin endpoint on the admin endpoint network.
- NetBackup for OpenStack appliance needs access to all endpoints of one type.

Recommendation: Provide access to all OpenStack Endpoint types

Cohesity recommends that you provide full access to all OpenStack endpoints to the NetBackup for OpenStack appliance to follow the OpenStack standards and best practices.

NetBackup for OpenStack generates its own endpoints as well. These endpoints point towards the NetBackup for OpenStack Appliance directly. This means that using those endpoints does not send the API calls towards the OpenStack Controller nodes first, but directly to the NetBackup for OpenStack virtual machine.

Following the OpenStack standards and best practices, it is therefore recommended to put the NetBackup for OpenStack endpoints on the same networks as the already existing OpenStack endpoints. This allows to extend the purpose of each endpoint type to the NetBackup for OpenStack service:

- The public endpoint to be used by OpenStack users when using NetBackup for OpenStack CLI or API.
- The internal endpoint to communicate with the OpenStack services.
- The admin endpoint to use the required admin only APIs of Keystone.

Backup target access required by NetBackup for OpenStack

The NetBackup for OpenStack solution uses backup target storage to securely place the backup data. NetBackup for OpenStack divides its backup data into two parts:

1. Metadata
2. Volume Disk Data

The first type of data is generated by the NetBackup for OpenStack appliance through communicating with the OpenStack Endpoints. All metadata that is stored together with a backup is written by the NetBackup for OpenStack Appliance to the backup target in the JSON format.

The second type of data is generated by the NetBackup for OpenStack datamover service running on the compute nodes. The nbosdm service reads the Volume Data

from the Cinder or Nova storage and transferring this data as QCOW2 image to the backup target. Each datamover service is hereby responsible for the virtual machines running on its compute node.

The network requirements are therefore:

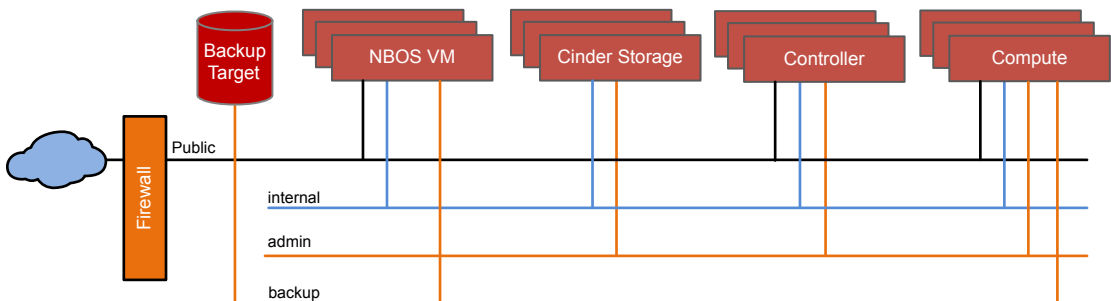
- The NetBackup for OpenStack appliance needs access to the backup target.
- Every compute node needs access to the backup target.

Example of a typical NetBackup for OpenStack network integration

Many OpenStack customers follow the OpenStack standards and best practices to have the public, internal, and admin endpoints on separate networks. They also typically don't have any network yet, which can access the desired backup target.

The starting network configuration typically looks as follows:

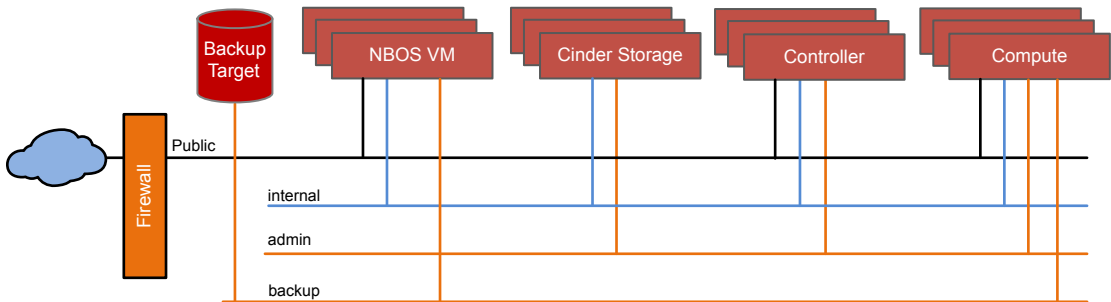
Figure 2-1 Typical OpenStack Network configuration before NetBackup for OpenStack gets installed



Following the OpenStack standards and Cohesity' recommendation the NetBackup for OpenStack Appliance is placed on all those three networks. Further is the access to the backup target that is required by NetBackup for OpenStack Appliance and Compute nodes. Here done by adding a 4th network.

The resulting network configuration looks as follows:

Figure 2-2 Typical OpenStack network configuration with NetBackup for OpenStack installed



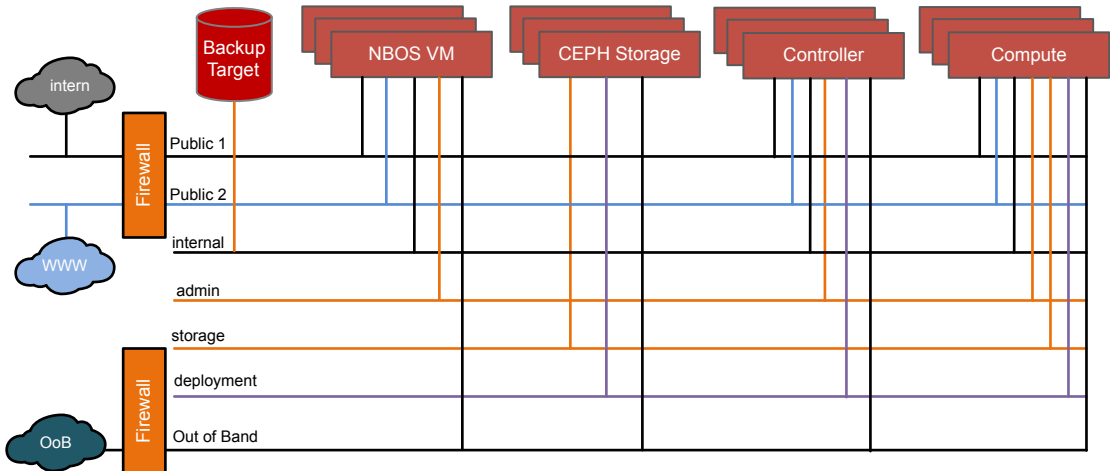
You can combine networks as necessary. As long as the required network access is available NetBackup for OpenStack works.

Other examples of NetBackup for OpenStack network integrations

Each OpenStack installation is different and so is the network configuration. There are endless possibilities of how to configure the OpenStack network and how to implement the NetBackup for OpenStack appliance into this network. The following three examples have been seen in production:

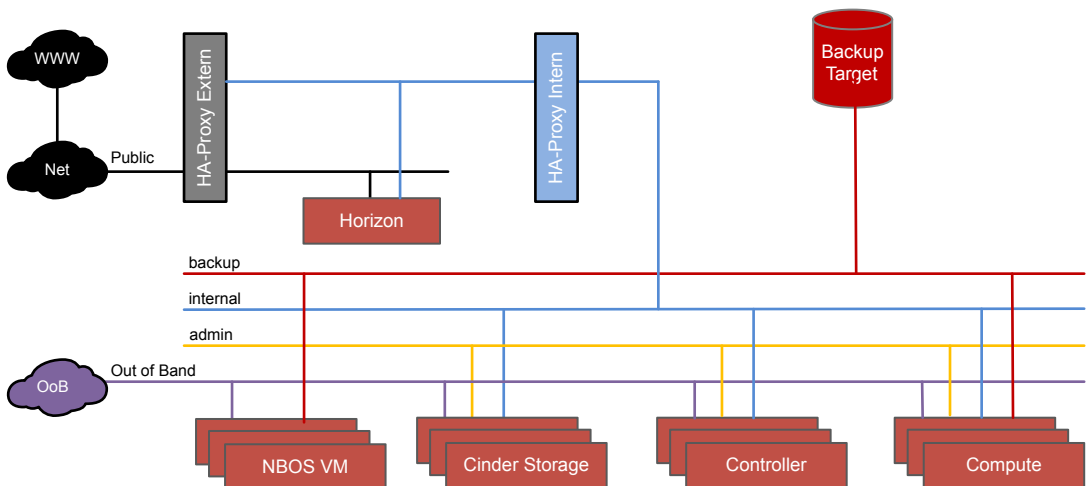
The first example is from a manufacturing company, which wanted to split the networks by function and decided to put the NetBackup for OpenStack backup target on the internal network as the backup and recovery function was identified as an OpenStack internal solution. This example looks complex but integrates NetBackup for OpenStack as recommended.

Figure 2-3 The split them all network example



The second example is from a financial institute that wanted to be sure that the OpenStack Users have no direct uncontrolled network access to the OpenStack infrastructure. Following this example requires additional work as the internal HA-Proxy needs to be configured to correctly translate the API calls towards the NetBackup for OpenStack

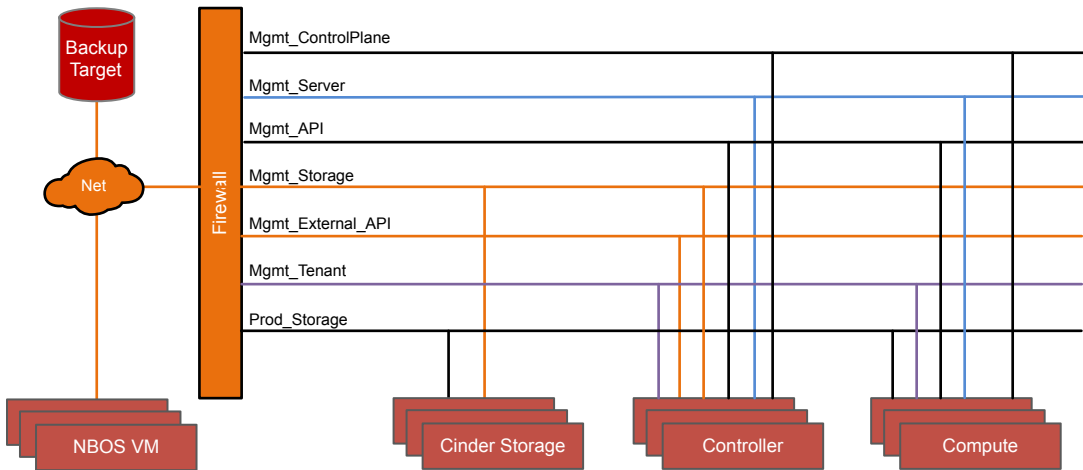
Figure 2-4 The no trust network example



The third example is from a service company that was forced to treat NetBackup for OpenStack as an external 3rd party solution, as we require a virtual machine

running outside of OpenStack. This kind of network configuration requires good planning on the NetBackup for OpenStack endpoints and firewall rules.

Figure 2-5 NetBackup for OpenStack as third party component network example



Preparing the installation

It is recommended to think about the following elements before the installation of NetBackup for OpenStack.

Tenant quotas

NetBackup for OpenStack uses Cinder snapshots for calculating full and incremental backups. For full backups, NetBackup for OpenStack creates Cinder snapshots for all the volumes in the backup job. It then leaves these Cinder snapshots behind for calculating the incremental backup image during the next backup. During an incremental backup operation it creates new Cinder snapshots, calculates the changed blocks between the new snapshots and the old snapshots that were left behind during the full/previous backups. It then deletes the old snapshots but leaves the newly created snapshots behind. So, it is important that each tenant that avails NetBackup for OpenStack backup functionality has sufficient Cinder snapshot quotas to accommodate these additional snapshots. The guideline is to add two snapshots for every volume that is added to backups to volume snapshot quotas for that tenant. You may also increase the volume quotas for the tenant by the same amount because NetBackup for OpenStack briefly creates a volume from a snapshot to read data from the snapshot for backup purposes. During a restore process,

NetBackup for OpenStack creates additional instances and Cinder volumes. To accommodate restore operations, a tenant should have a sufficient quota for Nova instances and Cinder volumes. Otherwise restore operations result in failures.

NetBackup for OpenStack Cluster

NetBackup for OpenStack can be deployed as a single node or a three-node cluster. We recommend that NetBackup for OpenStack is deployed as a three node cluster for fault tolerance and load balancing. NetBackup for OpenStack requires additional IP for cluster and is required for both single node and three node deployments. Cluster IP (virtual IP) is used to manage the cluster and is used to register NetBackup for OpenStack service endpoint in the keystone service catalog.

Spinning up the NetBackup for OpenStack virtual machine

The NetBackup for OpenStack Appliance is delivered as QCOW2 image and runs as a virtual machine on top of a KVM Hypervisor.

This guide shows the tested way to spin up the NetBackup for OpenStack Appliance on an RHV Cluster.

Creating the cloud-init image

The NetBackup for OpenStack appliance uses cloud-init to provide the initial network and user configuration.

Cloud-init reads its information from a metadata server or from a provided CD image. NetBackup for OpenStack uses the CD image.

Needed tools

To create the cloud-init image it is required to have genisoimage available.

```
#For RHEL
yum install genisoimage
```

Providing the metadata

Cloud-init uses two files for its metadata.

The first file is called `meta-data` and contains the information about the network configuration. The following is an example of this file.

```
[root@kvm]# cat meta-data
instance-id: NetBackup for OpenStack
network-interfaces: |
    auto ens3
    iface ens3 inet static
    address 158.69.170.20
    netmask 255.255.255.0
    gateway 158.69.170.30

    dns-nameservers 11.11.0.51
local-hostname: nbos-controller.domain.org
```

Warning: The instance-id has to match the virtual machine name in virsh.

The second file is called `user-data` and it contains little scripts and information for a setup. For example, the user passwords. The following is an example of this file.

```
[root@kvm]# cat user-data
#cloud-config
chpasswd:
  list: |
    root:password1
    stack:password2
  expire: False
```

Creating the image file

Both files metadata and user data is needed to create a working cloud-init image.

The image is created using `genisoimage` following this general command:

```
genisoimage -output <name>.iso -volid cidata -joliet -rock
</path/user-data> </path/meta-data>
```

An example of this command:

```
genisoimage -output nbos-firstboot-config.iso -volid cidata
-joliet -rock user-data meta-data
```


Spinning up the NetBackup for OpenStack appliance

After the cloud-init image is created, you can spin up the NetBackup for OpenStack appliance on the desired KVM server.

The following example command shows how to spin up the NetBackup for OpenStack appliance using the `virsh` command-line and the created ISO image.

```
virt-install -n nbosvm --memory 24576 --vcpus 8 \  
--os-type linux \  
--disk nbos-appliance-os-3.0.154.qcow2,device=disk,bus=virtio,size=40 \  
--network bridge=virbr0,model=virtio \  
--network bridge=virbr1,model=virtio \  
--graphics none \  
--import \  
--disk path=nbos-firstboot-config.iso,device=cdrom
```

For KVM server - RHEL version 9 and later:

```
virt-install -n nbosvm --memory 24576 --vcpus 8 \  
--osinfo rhel8-unknown \  
--disk nbos-appliance-os-3.0.154.qcow2,device=disk,bus=virtio,size=40 \  
--network bridge=virbr0,model=virtio \  
--network bridge=virbr1,model=virtio \  
--graphics none \  
--import \  
--disk path=nbos-firstboot-config.iso,device=cdrom
```

You can spin up the NetBackup for OpenStack appliance without a cloud-init iso-image. It spins up with default values.

Uninstalling cloud-init after first start

Once the NetBackup for OpenStack appliance is up and running with its initial configuration, it is recommended to uninstall cloud-init.

If cloud-init is not installed, it runs the network configuration again upon every start. Setting the network configuration back to DHCP, if no metadata is provided.

To uninstall cloud-init, follow the example below.

```
sudo yum remove cloud-init
```

Or

```
touch /etc/cloud/cloud-init.disabled
```

About NetBackup for OpenStack backup target types

NetBackup for OpenStack uses the universal share to store the backup images.

To configure the universal share, see the *Configuring and managing universal shares* chapter of the *NetBackup Deduplication Guide*.

While creating the universal share, in the **Host** field, add the IP addresses or the subnet of all the compute nodes and NetBackup for OpenStack virtual machines.

For example,

IP addresses: 10.210.xxx.xx1, 10.210.xxx.xx2, 10.210.xxx.xx3,
10.210.xxx.xx4, 10.210.xxx.xx5, 10.210.xxx.xx6, ... 10.210.xxx.x20

Subnet: 10.210.128.0/20

Note: Ensure that the compute nodes and NetBackup for OpenStack virtual machines are accessible from the media server to mount the universal share on the compute nodes and NetBackup for OpenStack virtual machines.

For the better performance, you can allocate the dedicated backup network to the compute nodes and use this network to mount the universal share.

Installing NetBackup for OpenStack Components

Once the NetBackup for OpenStack virtual machine or the cluster of NetBackup for OpenStack virtual machines are spun, the actual installation process can begin. This process contains the following steps:

1. Install the NetBackup for OpenStack datamover API (nbosdmapi) service on the control plane.
2. Install the NetBackup for OpenStack datamover (nbosdm) service on the compute plane.
3. Install the NetBackup for OpenStack Horizon plug-in into the Horizon service.

How these steps look in detail depends on the OpenStack distribution NetBackup for OpenStack is installed in. Each supported OpenStack distribution has its own deployment tools. NetBackup for OpenStack is integrated into these deployment tools to provide a native integration from the beginning to the end.

Installing on RHOSP

The Red Hat OpenStack Platform Director is the supported and recommended method to deploy and maintain any RHOSP installation.

NetBackup for OpenStack integrates natively into the RHOSP Director. Manual deployment methods are not supported for RHOSP.

Perform the following steps to install NetBackup for OpenStack on RHOSP.

Table 2-1 Installing on RHOSP

Step	Task	Description
1	Prepare for deployment.	See “Prepare for deployment” on page 27.
2	Upload NetBackup for OpenStack puppet module.	See “Uploading the NetBackup for OpenStack puppet module” on page 28.
3	Update overcloud roles data file to include NetBackup for OpenStack services.	See “Updating the overcloud roles data file to include NetBackup for OpenStack services” on page 28.
4	Prepare NetBackup for OpenStack container images..	See “Preparing the NetBackup for OpenStack container images” on page 29.
5	Provide environment details in nbos_env.yaml.	See “Providing the environment details in nbos_env.yaml” on page 30.
6	Deploy overcloud with NetBackup for OpenStack environment.	See “Deploying the overcloud with NetBackup OpenStack environment” on page 31.
7	Verify deployment.	See “Verifying the deployment” on page 32.
8	Perform additional steps on NetBackup for OpenStack Appliance.	See “Additional Steps on NetBackup for OpenStack Appliance” on page 34.
9	Troubleshoot for overcloud deployment failures.	See “Troubleshooting for overcloud deployment failures” on page 34.

Prepare for deployment

Perform the following tasks to prepare for the deployment:

- Select NetBackup for OpenStack backup target type.
See [“About NetBackup for OpenStack backup target types”](#) on page 26.
- Copy nbos-cfg-scripts to the undercloud.

See [“Copying nbos-cfg-scripts to the undercloud”](#) on page 28.

Copying nbos-cfg-scripts to the undercloud

Perform the following steps on the undercloud node on an already installed RHOSP environment. The overcloud-deploy command has to be run successfully already and the overcloud must be available.

Warning: All commands need to be run as user "stack" on undercloud node.

Run the following commands to copy the nbos-cfg-scripts:

```
cd /home/stack
cp <image location>/nbos-cfg-scripts.tar.gz /home/stack
gunzip /home/stack/nbos-cfg-scripts.tar.gz
tar xvf /home/stack/nbos-cfg-scripts.tar
cd nbos-cfg-scripts/redhat-director-scripts/<RHOSP_release_directory>/
```

Available RHOSP_release_directory values are:

- rhosp17.1

Uploading the NetBackup for OpenStack puppet module

The following commands upload the NetBackup for OpenStack puppet module to the overcloud registry. The actual upload happens upon the next deployment.

```
cd /home/stack/nbos-cfg-scripts/redhat-director-scripts/
<RHOSP_release_directory>/scripts/
./upload_puppet_module.sh
```

Updating the overcloud roles data file to include NetBackup for OpenStack services

NetBackup for OpenStack contains multiple services. Add these services to your roles_data.yaml.

If the roles_data.yaml is not customized, you can find it on the undercloud at the following location:

```
/usr/share/openstack-tripleo-heat-templates/roles_data.yaml
```

Add the following services to the roles_data.yaml.

Note: All commands need to be run as user "stack".

Adding the NetBackup for OpenStack datamover API service to role data file

This service needs to share the same role as the **keystone** and **database** service. In case of the predefined roles, these services run on the role **Controller**. In case of custom roles, it is necessary to use the same role where **OS::TripleO::Services::Keystone** service installed.

Add the following line to the identified role:

```
'OS::TripleO::Services::nbosdmapl'
```

Adding NetBackup for OpenStack datamover service to role data file

This service needs to share the same role as the nova-compute service. In case of the predefined roles, the nova-compute service runs on the role **Compute**. In case of custom defined roles, it is necessary to use the role that nova-compute service uses.

Add the following line to the identified role:

```
'OS::TripleO::Services::nbosdm'
```

Preparing the NetBackup for OpenStack container images

Warning: All commands need to be run as user "stack".

NetBackup for OpenStack uses the local registry on the undercloud to house packages.

NetBackup for OpenStack provides a shell script, which pushes the containers to the undercloud and updates the `nbos_env.yaml`.

```
cd
/home/stack/nbos-cfg-scripts/redhat-director-scripts/<RHOSP_release_directory>/scripts
sudo ./prepare_nbos_images.sh <UNDERCLOUD_REGISTRY_HOSTNAME>
<IMAGE_SOURCE_FOLDER>
```

Run following command to find UNDERCLOUD_REGISTRY_HOSTNAME.

In the following example `nbos-undercloud` is
`<UNDERCLOUD_REGISTRY_HOSTNAME>`

```
$ openstack tripleo container image list | grep keystone |
docker://nbos-undercloud:8787/rhosp-rhel9/openstack-keystone:17.1
```

CONTAINER_TAG format for RHOSP17.1: <NBOS_VERSION>-rhosp17.1

Example,

```
sudo ./prepare_nbos_images.sh nbos-undercloud 10.4.1.1035-rhosp17.1
/home/stack/nbos/nbos-rhosp17.1-10.4.1.1035
```

The changes can be verified using the following commands.

```
sudo podman images | grep nbos
localhost/nbos-horizon-plugin
10.4.1.1035-rhosp17.1 c4ba2c4ff0f8 3 days ago
1.01 GB
localhost/nbosdmapi
10.4.1.1035-rhosp17.1 8baac9920a8e 3 days ago
1.13 GB
localhost/nbosdm
10.4.1.1035-rhosp17.1 86542c17acc2 3 days ago
2.76 GB
```

```
(undercloud) [stack@host scripts]$ grep -i image
../environments/nbos_env.yaml
  docker_nbosdm_image:
nbos-undercloud:8787/nbosdm:10.4.1.1035-rhosp17.1
  docker_nbosdmapi_image:
nbos-undercloud:8787/nbosdmapi:10.4.1.1035-rhosp17.1
  ContainerHorizonImage:
nbos-undercloud:8787/nbos-horizon-plugin:10.4.1.1035-rhosp17.1
```

Providing the environment details in nbos_env.yaml

Provide the necessary details in the provided environment file. This environment file is used in the overcloud deployment to configure NetBackup for OpenStack components. Container image names have already been populated in the preparation of the container images. Still it is recommended to verify the container URLs.

The following information is required additionally:

- Network for the nbosdmapi
- nbosdm password

```
resource_registry:
  OS::TripleO::Services::nbosdm: ../services/nbosdm.yaml
  OS::TripleO::Services::nbosdmapi: ../services/nbosdmapi.yaml
```

```
# NOTE: If there are addition customizations to the endpoint map
(e.g. for
# other integrations), this will need to be regenerated.
OS::TripleO::EndpointMap: endpoint_map.yaml

parameter_defaults:

  ## Enable NetBackup for OpenStack's quota functionality on horizon
  ExtraConfig:
    horizon::customization_module: 'dashboards.overrides'

  ## Define network map for NetBackup OpenStack datamover API service
  ServiceNetMap:
    nbosdmapiNetwork: internal_api

  ## NetBackup for OpenStack datamover password for keystone and database
  nbosdmPassword: "test1234"

  ## NetBackup for OpenStack container pull urls
  docker_nbosdm_image: nbos-undercloud:8787/nbosdm:10.4.1.1035-rhosp17.1
  docker_nbosdmapi_image: nbos-undercloud:8787/nbosdmapi:10.4.1.1035-rhosp17.1

  ## If you do not want NetBackup for OpenStack's horizon plugin
  to replace your horizon container, just comment following line.
  ContainerHorizonImage: nbos-undercloud:8787/nbos-horizon-plugin:
  10.4.1.1035-rhosp17.1

  ## Don't edit following parameter
  EnablePackageInstall: True
```

Deploying the overcloud with NetBackup OpenStack environment

Use the following heat environment file and roles data file in overcloud deploy command:

1. nbos_env.yaml
2. roles_data.yaml
3. Use correct NetBackup OpenStack endpoint map file as per available Keystone endpoint configuration

Instead of `tls-endpoints-public-dns.yaml` file, use

`environments/nbos_env_tls_endpoints_public_dns.yaml`

Instead of `tls-endpoints-public-ip.yaml` file,

use `environments/nbos_env_tls_endpoints_public_ip.yaml`

Instead of `tls-everywhere-endpoints-dns.yaml` file,

use `environments/nbos_env_tls_everywhere_dns.yaml`

To include new environment files use `-e` option and for roles data file use `-r` option.

An example of overcloud deploy command:

```
openstack overcloud deploy --stack overcloud --templates \
  -n /home/stack/templates/network_data.yaml \
  -r /home/stack/templates/roles_data.yaml \
  -e /home/stack/templates/enable-tls.yaml \
  -e /home/stack/templates/inject-trust-anchor-hiera.yaml \
  -e /home/stack/nbos-cfg-scripts/redhat-director-scripts/
<RHOSP_RELEASE_DIRECTORY>/environments/nbos_env_tls_endpoints_public_ip.yaml
\
  -e /home/stack/templates/overcloud-baremetal-deployed.yaml \
  -e /home/stack/templates/overcloud-networks-deployed.yaml \
  -e /home/stack/templates/overcloud-vip-deployed.yaml \
  -e /home/stack/containers-prepare-parameter.yaml \
  -e /home/stack/templates/environment-file.yaml \
  -e /home/stack/nbos-cfg-scripts/redhat-director-scripts/
<RHOSP_release_directory>/environments/nbos_env.yaml \
  --ntp-server 172.16.8.24 >
/home/stack/templates/overcloud_deploy.log
```

Verifying the deployment

If the `nbosdmapi` container is not deployed on the controller node and the `nbosdm` container is not deployed on the compute node, perform the following steps:

1. Run the following command to render the templates with the modified `roles_data.yaml` file and perform the overcloud deployment.

```
/usr/share/openstack-tripleo-heat-templates/tools/process-templates.py
-p /usr/share/openstack-tripleo-heat-templates -r
/home/stack/templates/roles_data.yaml -n
/home/stack/templates/default-network-isolation.yaml -o
/home/stack/templates/generated-openstack-tripleo-heat-templates
--safe
```

2. Specify the generated template path with the `overcloud deploy` command.

For example,

```
openstack overcloud deploy --stack overcloud --templates
/home/stack/templates/generated-openstack-tripleo-heat-templates
```

If the containers are in restarting state or not listed by the following commands, your deployment is not done correctly.

- On the controller node:

Ensure that the NetBackup for OpenStack datamover API and horizon containers are in a running state and no other NetBackup for OpenStack container is deployed on controller nodes. When the role for these containers is not **controller**, check on the respective nodes according to configured `roles_data.yaml`.

```
[root@overcloud-controller-0 ~]# podman ps | grep nbos
26fcb9194566
rhospqa.ctlplane.localdomain:8787/nbosdmap:10.4.1.1035-rhosp17.1

kolla_start          5 days ago Up 5 days ago          nbosdmap
094971d0f5a9  rhospqa.ctlplane.localdomain:
8787/nbos-horizon-plugin:10.4.1.1035-rhosp17.1      kolla_start

5 days ago Up 5 days ago          horizon
```

- On the compute node:

Ensure that the NetBackup for OpenStack datamover API and horizon containers are in a running state and no other NetBackup for OpenStack container is deployed on controller nodes. When the role for these containers is not **controller**, check on the respective nodes according to configured `roles_data.yaml`.

```
[root@overcloud-controller-0 ~]# podman ps | grep nbos
26fcb9194566
rhospqa.ctlplane.localdomain:8787/nbosdmap:10.4.1.1035-rhosp17.1

kolla_start          5 days ago Up 5 days ago          nbosdmap
094971d0f5a9  rhospqa.ctlplane.localdomain:
8787/nbos-horizon-plugin:10.4.1.1035-rhosp17.1      kolla_start

5 days ago Up 5 days ago          horizon
```

- On the node with Horizon service:

Ensure that the horizon container is in the running state.

Note: The Horizon container is replaced with NetBackup for OpenStack Horizon container. This container has the latest OpenStack horizon + NetBackup for OpenStack horizon plug-in.

```
[root@overcloud-controller-0 ~]# podman ps | grep horizon
094971d0f5a9  rhospqa.ctlplane.localdomain:
8787/nbos-horizon-plugin:10.4.1.1035-rhosp17.1      kolla_start
5 days ago    Up 5 days ago          horizon
```

Additional Steps on NetBackup for OpenStack Appliance

Changing the nova user ID on the NetBackup for OpenStack Nodes

In RHOSP, "nova" user ID on nova-compute docker container is set to "42436". The "nova" user ID on the NetBackup for OpenStack nodes need to be set the same. Do the following steps on all NetBackup for OpenStack nodes:

1. Execute the script.
2. Verify that nova user and group ID has changed to 42436.

```
## Execute the shell script to change 'nova' user and group id to '42436'
$ ./home/stack/nova_userid.sh

## Ignore any errors and verify that 'nova' user and group id has
changed to '42436'
$ id nova
uid=42436(nova) gid=42436(nova) groups=42436(nova),990(libvirt),36(kvm)
```

Troubleshooting for overcloud deployment failures

NetBackup for OpenStack components are deployed using the puppet scripts.

If the overcloud deployment fails, run the following command to get the list of errors.

- (undercloud)\$ openstack stack failures list <overcloud>--long
<overcloud> The name of the overcloud.
- (undercloud)\$ openstack stack list --nested --property
status=FAILED

For more information, see the *OpenStack documentation*.

If the nbosdmapi container does not start or is in the restarting state, run the following commands to get the logs to troubleshoot.

- podman logs nbosdmapi

- `tail -f /var/log/containers/nbosdmap1/nbosdmap1.log`

If the nbosdm container does not start or is in the restarting state, run the following commands to get the logs to troubleshoot.

- `podman logs nbosdm`
- `tail -f /var/log/containers/nbosdm/nbosdm.log`

Installing on Ansible OpenStack Ussuri

Perform the following steps to install NetBackup for OpenStack on Ansible OpenStack Ussuri

Table 2-2 Installing on Ansible OpenStack Ussuri

Step	Task	Description
1	Verify that file-level logging is configured for OpenStack components on Horizon container	See “Verify that file-level logging is configured for OpenStack components on Horizon container” on page 35.
2	Change the nova user ID on the NetBackup for OpenStack Nodes	See “Changing the nova user ID on the NetBackup for OpenStack Nodes” on page 36.
3	Prepare deployment host	See “Preparing the deployment host” on page 37.
4	Deploy NetBackup for OpenStack components	See “Deploying the NetBackup for OpenStack components” on page 40.
5	Verify the NetBackup for OpenStack deployment	See “Verifying the NetBackup for OpenStack deployment” on page 41.

Verify that file-level logging is configured for OpenStack components on Horizon container

NetBackup for OpenStack Horizon plug-in uses OpenStack’s logging services to store the logs. It is recommended that you configure system logging for OpenStack components on Horizon container.

Ensure that you configure the following parts of the logging to generate structured log information to a file.

Sample configuration:

- **Formatters:** Define the formatting of log information in the log file.

```
'verbose': {
    'format': '%(asctime)s %(process)d %(levelname)s %(name)s %(message)s'
},
```

- **Handlers:** Add a file handler to write log information to the log file.

```
'file': {
    'level': 'DEBUG',
    'class': 'logging.FileHandler',
    'filename': '/var/log/horizon/horizon.log',
    'formatter': 'verbose',
},
```

- **Loggers:** Update each OpenStack component in use with the file handler information to the log file.

For example, OpenStack dashboard, Horizon, Nova client, Cinder client, Keystone client, Glance client, Neutron client, OpenStack authorization, Django, and so on.

```
'horizon': {
    'handlers': ['file'],
    'level': 'DEBUG',
    'propagate': False,
}
```

It is recommended that you enable log rotation to restrict the volume of the log data to avoid overflowing the record store. For more information about logging and configuring log rotation, see *Django documentation*.

Changing the nova user ID on the NetBackup for OpenStack Nodes

NetBackup for OpenStack virtual machine uses the nova user ID and group ID 162:162 by default. Ansible OpenStack is not always nova user ID 162 on nova-compute containers. The nova user ID on the NetBackup for OpenStack virtual machine nodes must be same as the nova-compute containers. If nova ID is not 162:162, perform the following steps on all NetBackup for OpenStack virtual machine nodes.

Before you perform the following steps, verify that the user ID and group ID is not used by any other services on NetBackup for OpenStack virtual machine. For example, If nova ID on compute node is 997, verify that user ID is not used by any other services on NetBackup for OpenStack virtual machine. If 997 user ID is

assigned to `rabbitmq` and 997 group ID is assigned to `SSH` service on NetBackup for OpenStack virtual machine, you must free this ID.

```
#cat /etc/passwd | grep 997
#pid 997
#ps -ef | grep 997
#usermod -u 900 rabbitmq
#cat /etc/group | grep 997
#groupmod -g 901 ssh_keys
#reboot
```

1. Go to the directory `/home/stack .`
2. Assign the executable permissions to `nova_userid.sh` file.

```
#chmod +x nova_userid.sh
```

3. Edit script to use the correct nova ID.
4. Execute the script.

```
#./nova_userid.sh
```

5. Verify that nova user and group ID has changed to the desired value.

```
#id nova
```

Preparing the deployment host

Select the NetBackup for OpenStack backup target storage type.

See See [“About NetBackup for OpenStack backup target types”](#) on page 26.

Copy Ansible roles and vars to the required places.

```
cd nbos-cfg-scripts/
cp -R ansible/roles/* /opt/openstack-ansible/playbooks/roles/
cp ansible/main-install.yml /opt/openstack-ansible/playbooks/
os-nbos-install.yml
cp ansible/environments/group_vars/all/vars.yml /etc/openstack_
deploy/user_nbos_vars.yml
```

Add NetBackup for OpenStack playbook to

`/opt/openstack-ansible/playbooks/setup-openstack.yml` at the end of the file.

```
- import_playbook: os-nbos-install.yml
```

Add the following information at the end of the file

/etc/openstack_deploy/user_variables.yml

```
# Datamover haproxy setting
haproxy_extra_services:
- service:
    haproxy_service_name: nbosdm_service
    haproxy_backend_nodes: "{{ groups['nbosdmapi_all'] | default([]) }}"
    haproxy_ssl: "{{ haproxy_ssl }}"
    haproxy_port: 8784
    haproxy_balance_type: http
    haproxy_balance_alg: roundrobin
    haproxy_timeout_client: 10m
    haproxy_timeout_server: 10m
    haproxy_backend_options:
      - "httpchk GET / HTTP/1.0\r\nUser-agent:\\ osa-haproxy-healthcheck"
```

Create the file /opt/openstack-ansible/inventory/env.d/nbos-nbosdmapi.yml

Add the following information to the file.

```
cat > /opt/openstack-ansible/inventory/env.d/nbos-nbosdmapi.yml
```

```
component_skel:
  nbosdmapi_api:
    belongs_to:
      - nbosdmapi_all

container_skel:
  nbosdmapi_container:
    belongs_to:
      - nbos-nbosdmapi_containers
  contains:
    - nbosdmapi_api

physical_skel:
  nbos-nbosdmapi_containers:
    belongs_to:
      - all_containers
  nbos-nbosdmapi_hosts:
    belongs_to:
      - hosts
```

Edit the file `/etc/openstack_deploy/openstack_user_config.yml` according to the example below to set host entries for NetBackup for OpenStack components.

```
#nbosdmapi
nbos-nbosdmapi_hosts:      # Add controller details in this section as
                           # nbos-dmapi is resides on controller nodes.

    infra1:                # Controller host name
        ip: <controller_ip> # IP address of controller
    infra2:                # For multiple controller nodes add controller node
                           # details in same manner as shown in infra2
        ip: <controller_ip>

#nbos-datamover
nbos_compute_hosts:       # Add compute details in this section as nbosdm
                           # resides on compute nodes.

    infra-1:               # Compute host name
        ip: <compute_ip>   # IP address of compute
    infra2:                # For multiple compute nodes add compute node
                           # details in same manner as shown in infra2
        ip: <compute_ip>
```

Edit the common editable parameter section in the file

`/etc/openstack_deploy/user_nbos_vars.yml`

Append the required details like NetBackup for OpenStack Appliance IP address, NetBackup for OpenStack package version, OpenStack distribution, snapshot storage backend, SSL related information and so on.

```
##common editable parameters required for installing nbos-horizon-plugin,
nbosdm and nbosdmapi
#ip address of nbosvm
IP_ADDRESS: <Nbosvm IP>
##Time Zone
TIME_ZONE: "Etc/UTC"

#Update NBOS package version here, we will install mentioned version
plugins for Example# NBOS_PACKAGE_VERSION: 3.3.36
NBOS_PACKAGE_VERSION: <Build No>
# Update Openstack dist code name like ussuri etc.
OPENSTACK_DIST: ussuri

#Need to add the following statement in nova sudoers file
```

```
#nova ALL = (root) NOPASSWD: /home/nbos/.virtenv/bin/privsep-helper *
#These changes require for nbosdm, Otherwise nbosdm will not work
#Are you sure? Please set variable to
# UPDATE_NOVA_SUDOERS_FILE: proceed
#other wise ansible nbosdm installation will exit
UPDATE_NOVA_SUDOERS_FILE: proceed

###details of nbosdmapl
##If SSL is enabled "NBOSDMAPI_ENABLED_SSL_APIS" value should be nbosdmapl.
#NBOSDMAPI_ENABLED_SSL_APIS: nbosdmapl
##If SSL is disabled "NBOSDMAPI_ENABLED_SSL_APIS" value should be empty.
NBOSDMAPI_ENABLED_SSL_APIS: ""
NBOSDMAPI_SSL_CERT: ""
NBOSDMAPI_SSL_KEY: ""

#### Any service is using Ceph Backend then set ceph_backend_enabled
value to True
#True/False
ceph_backend_enabled: False

#Set verbosity level and run playbooks with -vvv option to display
custom debug messages
verbosity_level: 3
```

Deploying the NetBackup for OpenStack components

Run the following commands to deploy only NetBackup for OpenStack components in case of an already deployed Ansible OpenStack.

```
cd /opt/openstack-ansible/playbooks

# To create nbosdmapl container
openstack-ansible lxc-containers-create.yml

#To Deploy NetBackup for OpenStack components
openstack-ansible os-nbos-install.yml

#To configure Haproxy for nbosdmapl
openstack-ansible haproxy-install.yml
```


If Ansible OpenStack is not already deployed, run the native OpenStack deployment commands to deploy OpenStack and NetBackup for OpenStack components together. An example for the native deployment command is given below:

```
openstack-ansible setup-infrastructure.yml --syntax-check
openstack-ansible setup-hosts.yml
openstack-ansible setup-infrastructure.yml
openstack-ansible setup-openstack.yml
```

Verifying the NetBackup for OpenStack deployment

Verify that the NetBackup for OpenStack datamover API service is deployed and has started. Run the following commands on controller node.

```
lxc-ls # Check the nbosdmapi container is present on controller node.
lxc-info -s controller_nbosdmapi_container-all984bf
# Confirm running status of the container
```

Verify that the NetBackup for OpenStack datamover service is deployed and has started on compute nodes. Run the following command on compute nodes.

```
systemctl status nbosdm.service
```

Verify that the NetBackup for OpenStack horizon plugin, nbosdmclient, and nbosjmcclient are installed on the Horizon container.

Run the following command on Horizon container.

```
lxc-attach -n controller_horizon_container-1d9c055c
# To login on horizon container
apt list | egrep 'nbos-horizon-plugin|nbosjmcclient|nbosdmclient '
# For ubuntu based container
yum list installed |egrep 'nbos-horizon-plugin|nbosjmcclient|
nbosdmclient '
# For RHEL based container
```

Run the following commands to verify haproxy setting on controller node.

```
haproxy -c -V -f /etc/haproxy/haproxy.cfg # Verify the keyword
nbosdm_service-back is present in output.
```

Installing on Kolla

Perform the following steps to install NetBackup for OpenStack on Kolla.

Table 2-3 Installing on Kolla

Step	Task	Description
1	Changing the nova user ID on the NetBackup for OpenStack Nodes	See “Changing the nova user ID on the NetBackup for OpenStack Nodes” on page 42.
2	Select backup target type.	See “About NetBackup for OpenStack backup target types” on page 26.
3	Copy the NetBackup for OpenStack deployment scripts	See “Copying the NetBackup for OpenStack deployment scripts” on page 43.
4	Copy the NetBackup for OpenStack deployment scripts to Kolla-ansible deploy scripts	See “Copying the NetBackup for OpenStack deployment scripts to Kolla-ansible deploy scripts” on page 44.
5	Push the NetBackup for OpenStack images to the local registry	See “Pushing NetBackup for OpenStack images to the local registry” on page 45.
6	Edit globals.yml to set NetBackup for OpenStack parameters	See “Editing globals.yml to set NetBackup for OpenStack parameters” on page 51.
7	Enable NetBackup for OpenStack backup mount feature	See “Enabling the NetBackup for OpenStack backup mount feature” on page 52.
7	Pull NetBackup for OpenStack container images	See “Pulling the NetBackup for OpenStack container images” on page 54.
8	Deploy NetBackup for OpenStack components	See “Deploying the NetBackup for OpenStack components” on page 54.
9	Verify NetBackup for OpenStack deployment	See “Verifying the NetBackup for OpenStack deployment” on page 54.

Changing the nova user ID on the NetBackup for OpenStack Nodes

NetBackup for OpenStack virtual machine uses the nova user ID and group ID 162:162 by default. Kolla OpenStack is not always nova user ID 162 on nova-compute containers. The nova user ID on the NetBackup for OpenStack virtual

machine nodes must be same as the nova-compute containers. If nova ID is not 162:162, perform the following steps on all NetBackup for OpenStack virtual machine nodes.

Before you perform the following steps, verify that the user ID and group ID is not used by any other services on NetBackup for OpenStack virtual machine. For example, If nova ID on compute node is 997, verify that user ID is not used by any other services on NetBackup for OpenStack virtual machine. If 997 user ID is assigned to `rabbitmq` and 997 group ID is assigned to `ssh` service on NetBackup for OpenStack virtual machine, you must free this ID.

```
#cat /etc/passwd | grep 997
#pid 997
#ps -ef | grep 997
#usermod -u 900 rabbitmq
#cat /etc/group | grep 997
#groupmod -g 901 ssh_keys
#reboot
```

1. Go to the directory `/home/stack` .
2. Assign the executable permissions to `nova_userid.sh` file.

```
#chmod +x nova_userid.sh
```

3. Edit script to use the correct nova ID.
4. Execute the script.

```
#./nova_userid.sh
```

5. Verify that nova user and group ID has changed to the desired value.

```
#id nova
```

Copying the NetBackup for OpenStack deployment scripts

To copy the NetBackup for OpenStack deployment scripts

1. Ensure that the nbos-cfg-scripts are available on Kolla ansible server at `/root` or any other directory.
2. Create and switch to directory to untar the NetBackup for OpenStack deployment scripts.

```
mkdir nbos-cfg-scripts
cd nbos-cfg-scripts/
```

3 Untar the tar file.

```
tar -xvf nbos-cfg-scripts-<NBOS version number>.tar.gz
```

For example, `tar -xvf nbos-cfg-scripts-9.1.2.20211021104525.tar.gz`

4 Copy NetBackup for OpenStack ansible role into Kolla-ansible roles directory.

```
cp -R kolla/roles/NetBackupOpenStack  
/path/to/venv/share/kolla-ansible/ansible/roles/
```

Copying the NetBackup for OpenStack deployment scripts to Kolla-ansible deploy scripts

To copy the NetBackup for OpenStack deployment scripts to Kolla-ansible deploy scripts

1 Go to the installation directory.

Yoga release: `cd kolla_yoga`

Caracal release: `cd kolla`

2 Add NetBackup for OpenStack global variables to `globals.yml`.

Take backup of `globals.yml`.

```
cp /etc/kolla/globals.yml /opt/
```

Append NetBackup for OpenStack global variables to `globals.yml`.

```
cat NetBackupOpenStack_globals.yml >> /etc/kolla/globals.yml
```

3 Add NetBackup for OpenStack passwords to `kolla passwords.yml` file.

Append `NetBackupOpenStack_passwords.yml` to `/etc/kolla/passwords.yml`.

Passwords are empty. Set these passwords manually in the `/etc/kolla/passwords.yml`.

Take backup of `passwords.yml`.

```
cp /etc/kolla/passwords.yml /opt/
```

Append NetBackup for OpenStack global variables to `passwords.yml`.

```
cat NetBackupOpenStack_passwords.yml >> /etc/kolla/passwords.yml
```

Edit `/etc/kolla/passwords.yml`. At the end of the file, set NetBackup for OpenStack passwords.

```
NetBackupOpenStack_keystone_password: <password>
```

```
NetBackupOpenStack_database_password: <password>
```

- 4 Append the NetBackup for OpenStack's YAML file content to the kolla ansible's site.yml file.

Note: Before you append the YAML file content, take the backup of the site.yml file.

```
cp /path/to/venv/share/kolla-ansible/ansible/site.yml /opt/
```

Caracal release: `cat NetBackupOpenStack_site.yml >> /path/to/venv/share/kolla-ansible/ansible/site.yml`

Yoga release: `cat NetBackupOpenStack_yoga_site.yml >> /path/to/venv/share/kolla-ansible/ansible/site.yml`

- 5 Append NetBackupOpenStack_inventory.txt to your cloud's kolla-ansible inventory file.

```
cat NetBackupOpenStack_inventory.txt >> <your inventory file name path>
```

For example,

```
cat NetBackupOpenStack_inventory.txt >> /root/multinode
```

Pushing NetBackup for OpenStack images to the local registry

Perform the following tasks to push NetBackup for OpenStack images to local registry.

Table 2-4

Step	Task	Description
1	Run the local registry.	See “Running the local registry” on page 45.
2	Load the images from tar and push them to the local repository	See “Loading the images from tar and pushing them to the local repository” on page 46.

Running the local registry

Run the local registry to get container images of NetBackup for OpenStack on RHEL and Ubuntu.

To run the local registry

- ◆ Run the following command on the deployment node to start the registry container.

```
docker run -d -p 5001:5000 --restart=always --name
<local_registry_name> registry:2
```

<local_registry_name> Your registry name. If your registry name does not have a name, assign a new name. The command pulls the registry image from docker.io and runs that container.

Loading the images from tar and pushing them to the local repository

Ensure that the proper tar files of nbosdmapi, nbosdm and nbos-horizon-plugin are available on the deployment node.

NBOS_Version	NetBackup for OpenStack version number.
kolla-base-distro	Ubuntu
kolla-install-type	Binary or source
FQDN	Hostname of kolla deployment server.

To load the images from tar and push them to the local repository

- 1 Load NetBackup for OpenStack images from the tar file.

Run the following commands:

- nbosdmapi

```
docker load --input nbosdmapi-{{ kolla-base-distro }}:{{
NBOS_version }}-{{ openstack_release }}.tar
```

For example,

```
docker load --input nbosdmapi-ubuntu-9.1.2.20211021104525-{{
openstack_release }}.tar
```

- nbosdm

```
docker load --input nbosdm-{{ kolla-base-distro }}:{{
NBOS_version }}-{{ openstack_release }}.tar
```

For example,

```
docker load --input nbosdm-ubuntu-9.1.2.20211021104525-{{
openstack_release }}.tar
```

- nbos-horizon-plugin (Yoga release)

```
docker load --input nbos-horizon-plugin-{{ kolla-install-type
}}-{{ kolla-base-distro }}:{{ NBOS_version }}-{{
openstack_release }}.tar
```

For example,

```
docker load --input
nbos-horizon-plugin-source-ubuntu-9.1.2.20211021104525-{{
openstack_release }}.tar
```

- **nbos-horizon-plugin (Caracal release)**

```
docker load --input nbos-horizon-plugin-{{ kolla-base-distro
}}:{{ NBOS_version }}-{{ openstack_release }}.tar
```

For example,

```
docker load --input
nbos-horizon-plugin-ubuntu-9.1.2.20211021104525-{{
openstack_release }}.tar
```

2 Tag the NetBackup for OpenStack images with appropriate name.

Run the following commands:

- **nbosdmapi**

- `docker tag localhost/nbosdmapi-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }} nbos/nbosdmapi-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`

- `docker tag localhost/nbosdmapi-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }} FQDN:5001/nbos/nbosdmapi-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`

Examples,

- `docker tag localhost/nbosdmapi-ubuntu:9.1.2.20211021104525-{{ openstack_release }} nbos/nbosdmapi-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`

- `docker tag localhost/nbosdmapi-ubuntu:9.1.2.20211021104525-{{ openstack_release }} deployment-vm.vxindia.veritas.com:5001/nbos/nbosdmapi-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`

- **nbosdm**

- `docker tag localhost/nbosdm-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`
`nbos/nbosdm-<kolla-base-distro>:<NBOS_version>-{{ openstack_release }}`
- `docker tag localhost/nbosdm-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`
`FQDN:5001/nbos/nbosdm-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`

Examples,

- `docker tag localhost/nbosdm-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`
`nbos/nbosdm-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`
- `docker tag localhost/nbosdm-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`
`deployment-vm.vxindia.veritas.com:5001/nbos/nbosdm-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`
- **nbos-horizon-plugin (Yoga release)**
 - `docker tag localhost/nbos-horizon-plugin-{{ kolla-install-type }}-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`
`nbos/nbos-horion-plugin-{{ kolla-install-type }}-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`
 - `docker tag localhost/nbos-horizon-plugin-{{ kolla-install-type }}-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`
`FQDN:5001/nbos/nbos-horizon-plugin-{{ kolla-install-type }}-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`

Examples,

- `docker tag localhost/nbos-horizon-plugin-source-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`
`nbos/nbos-horizon-plugin-source-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`
- `docker tag localhost/nbos-horizon-plugin-source-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`
`deployment-vm.vxindia.veritas.com:5001/`


```
nbos/nbos-horizon-plugin-source-ubuntu:9.1.2.20211021104525-{{
  openstack_release }}
```

- **nbos-horizon-plugin (Caracal release)**

- `docker tag localhost/nbos-horizon-plugin-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`
`nbos/nbos-horizon-plugin-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`
- `docker tag localhost/nbos-horizon-plugin-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`
`FQDN:5001/nbos/nbos-horizon-plugin-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`

Examples,

- `docker tag`
`localhost/nbos-horizon-plugin-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`
`nbos/nbos-horizon-plugin-source-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`
- `docker tag`
`localhost/nbos-horizon-plugin-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`
`deployment-vm.vxindia.veritas.com:5001/nbos/nbos-horizon-plugin-source-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`

3 Push the tagged image to local registry.

Run the following commands:

- **nbosdmapi**
`docker push FQDN:5001/nbos/nbosdmapi-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`
 For example,
`docker push`
`deployment-vm.vxindia.veritas.com:5001/nbos/nbosdmapi-ubuntu:9.1.2.20211021104525-{{ openstack_release }}`
- **nbosdm**
`docker push FQDN:5001/nbos/nbosdm-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}`
 For example,

```
docker push deployment-vm.vxindia.veritas.com:5001/nbos/nbosdm-ubuntu:9.1.2.20211021104525-{{ openstack_release }}
```

- **nbos-horizon-plugin (Yoga release)**

```
docker push FQDN:5001/nbos/nbos-horizon-plugin-{{ kolla-install-type }}-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}
```

For example,

```
docker push deployment-vm.vxindia.veritas.com:5001/nbos/nbos-horizon-plugin-source-ubuntu:9.1.2.20211021104525-{{ openstack_release }}
```

- **nbos-horizon-plugin (Caracal release)**

```
docker push FQDN:5001/nbos/nbos-horizon-plugin-{{ kolla-base-distro }}:{{ NBOS_version }}-{{ openstack_release }}
```

For example,

```
docker push deployment-vm.vxindia.veritas.com:5001/nbos/nbos-horizon-plugin-ubuntu:9.1.2.20211021104525-{{ openstack_release }}
```

4 Add the `insecure-registries` entry in `/etc/docker/daemon.json` on all controller and compute nodes.

Open the `daemon.json` file and make the changes as follows:

```
cat /etc/docker/daemon.json
{
  "log-opts": {
    "max-file": "5",
    "max-size": "50m"
  },
  "registry-mirrors": [
    "http://<deployment node ip>:4000"
  ],
  "insecure-registries": [
    "FQDN:5001"
  ]
}
```

- 5** Add the `insecure-registries` entry in `/etc/docker/daemon.json` on deployment nodes.

If `/etc/docker/` directory does not exist, create it and create `daemon.json` file.

Open the `daemon.json` file and make the changes as follows:

```
cat /etc/docker/daemon.json
{ "insecure-registries":["FQDN:5001"] }
```

- 6** Restart the docker.

```
systemctl restart docker
```

- 7** Verify that the specified images are pushed in the registry.

- Controller and compute nodes: `curl -X GET http://FQDN:5001/v2/_catalog`
- Deployment node: `docker info`

For example,

```
curl -X GET http://deployment-vm.vxindia.veritas.com:5001/v2/_catalog
```

Sample output:

```
curl -X GET http://deployment-vm.vxindia.veritas.com:5001/v2/_catalog
//Output should look like below:
{"repositories":["nbos/nbos-horizon-plugin-source-ubuntu",
"nbos/nbosdm-ubuntu", "nbos/nbosdmapi-ubuntu"]}
```

Editing `globals.yml` to set NetBackup for OpenStack parameters

Edit `/etc/kolla/globals.yml` file to configure NetBackup for OpenStack backup target and build details. You can find the NetBackup for OpenStack related parameters at the end of `globals.yml` file. You must configure the information such as NetBackupOpenStack tag, backup target type, and backup target details.

Following is the list of parameters that you can edit.

Table 2-5 globals.yml parameters

Parameter	Description
NetBackupOpenStack_tag	Container tags.
horizon_image_full	By default, the NetBackup for OpenStack Horizon container does not get deployed. Uncomment this parameter to deploy NetBackup for OpenStack Horizon container instead of Openstack Horizon container.
NetBackupOpenStack_docker_username	Default docker user of NetBackup for OpenStack. (read permission only)
NetBackupOpenStack_docker_password	Password for default docker user of NetBackup for OpenStack.
NetBackupOpenStack_docker_registry	Local registry name, which contains NetBackup for OpenStack component images.

Enabling the NetBackup for OpenStack backup mount feature

To enable NetBackup for OpenStack's backup mount feature it is necessary to make the NetBackup for OpenStack Backup target available to the nova-compute and nova-libvirt containers.

Edit

/path/to/venv/share/kolla-ansible/ansible/roles/nova-cell/defaults/main.yml and find `nova_libvirt_default_volumes` variable. Append the NetBackup for OpenStack mount bind `/var/nbos:/var/nbos:shared` to the list of already existing volumes.

For a default Kolla installation, the variables look as follows:

```
nova_libvirt_default_volumes:
  - "{{ node_config_directory }}/nova-libvirt/{{ container_config_directory }}:ro"
  - "/etc/localtime:/etc/localtime:ro"
  - "{{ '/etc/timezone:/etc/timezone:ro' if ansible_os_family == 'Debian' else '' }}"
  - "/lib/modules:/lib/modules:ro"
  - "/run/:/run/:shared"
  - "/dev:/dev"
```

```

- "/sys/fs/cgroup:/sys/fs/cgroup"
- "kolla_logs:/var/log/kolla/"
- "libvirtd:/var/lib/libvirt"
- "{{ nova_instance_datadir_volume }}:/var/lib/nova/"
- "{% if enable_shared_var_lib_nova_mnt | bool %}/var/lib/nova/mnt:/var/lib/nova/mnt:shared{% endif %}"
- "nova_libvirt_qemu:/etc/libvirt/qemu"
- "{{ kolla_dev_repos_directory ~ '/nova/nova:/var/lib/kolla/venv/lib/python' ~ distro_python_version ~ '/site-packages/nova' if nova_dev_mode | bool else '' }}"
- "/var/nbos:/var/nbos:shared"

```

Next, find the variable `nova_compute_default_volumes` in the same file and append the mount bind `/var/nbos:/var/nbos:shared` to the list.

After the change, the variable will look as follows for a default Kolla installation :

```

nova_compute_default_volumes:
- "{{ node_config_directory }}/nova-compute:{{ container_config_directory }}:/ro"
- "/etc/localtime:/etc/localtime:ro"
- "{{ '/etc/timezone:/etc/timezone:ro' if ansible_os_family == 'Debian' else '' }}"
- "/lib/modules:/lib/modules:ro"
- "/run:/run:shared"
- "/dev:/dev"
- "kolla_logs:/var/log/kolla/"
- "{% if enable_iscsid | bool %}iscsi_info:/etc/iscsi{% endif %}"
- "libvirtd:/var/lib/libvirt"
- "{{ nova_instance_datadir_volume }}:/var/lib/nova/"
- "{% if enable_shared_var_lib_nova_mnt | bool %}/var/lib/nova/mnt:/var/lib/nova/mnt:shared{% endif %}"
- "{{ kolla_dev_repos_directory ~ '/nova/nova:/var/lib/kolla/venv/lib/python' ~ distro_python_version ~ '/site-packages/nova' if nova_dev_mode | bool else '' }}"
- "/var/nbos:/var/nbos:shared"

```

In case of using Ironic compute nodes one more entry need to be adjusted in the same file. Find the variable `nova_compute_ironic_default_volumes` and append NBOS mount `/var/nbos:/var/nbos:shared` to the list.

After the changes the variable will looks like the following:

```
nova_compute_ironic_default_volumes:
- "{{{ node_config_directory }}/nova-compute-ironic/{{{ container_config_
  directory }}}:ro"
- "/etc/localtime:/etc/localtime:ro"
- "{{{ '/etc/timezone:/etc/timezone:ro' if ansible_os_family == 'Debian'
  else '' }}}"
- "kolla_logs:/var/log/kolla/"
- "{{{ kolla_dev_repos_directory ~ '/nova/nova:/var/lib/kolla/venv/lib/
  python' ~ distro_python_version ~ '/site-packages/nova' if nova_dev
  _mode | bool else '' }}}"
- "/var/nbos:/var/nbos:shared"
```

Pulling the NetBackup for OpenStack container images

Pull the NetBackup for OpenStack container images from the dockerhub based on the existing inventory file.

```
kolla-ansible -i <inventory file name> pull --tags NetBackup for
OpenStack
```

For example,

```
kolla-ansible -i multinode pull --tags netbackup
```

Deploying the NetBackup for OpenStack components

Run the following deployment command using the existing inventory file.

```
kolla-ansible -i <inventory file name> deploy
```

For example,

```
kolla-ansible -i multinode deploy
```

Verifying the NetBackup for OpenStack deployment

Verify that the nodes that run the NetBackup for OpenStack containers are available and healthy.

```
docker ps | grep nbosdmapi
```

Sample output:

```
3107046dce84    r0000-000-vm00.sample.name.com:
5001/nbos/nbosdmapi-ubuntu:10.0.0.1.1007-{{{ openstack_release }}}
"dumb-init --single-..."    9 days ago    Up 9 days
NetBackupOpenStack_datamover_api
```

```
docker ps | grep nbosdm
```

Sample output:

```
77f22039bd54    r0000-000-vm00.sample.name.com:
5001/nbos/nbosdm-ubuntu:10.0.0.1.1007-{{ openstack_release }}
"dumb-init
--single-..."    9 days ago      Up 4 days
NetBackupOpenStack_datamover
```

```
docker ps | grep horizon
```

Sample output:

```
dde1c91ed1a0    r0000-000-vm00.sample.name.com:
5001/nbos/nbos-horizon-plugin-binary-ubuntu:10.0.0.1.1007-{{
openstack_release }}
"dumb-init --single-..."    7 months ago    Up 7 months
horizon
```

Configuring NetBackup for OpenStack

NetBackup for OpenStack configuration process uses Ansible scripts. Ansible, in the last few years, has grown in popularity as a preferred configuration management tool and NetBackup for OpenStack uses Ansible playbooks extensively to configure the NetBackup for OpenStack cluster. To troubleshoot NetBackup for OpenStack configuration issues, the user should have a basic understanding of Ansible playbook output.

Ansible modules are inherently idempotent and hence NetBackup for OpenStack configuration can run any number of times to change or reconfigure NetBackup for OpenStack cluster.

Once the virtual machine is started, point your browser (Chrome or Firefox) to NetBackup for OpenStack node IP address.

This brings you to the NetBackup for OpenStack dashboard, which contains the NetBackup for OpenStack configurator.

The user is: admin The default password is: password

After the very first login, you are requested to change the admin password.

NetBackup for OpenStack requires you to configure the cluster once and the NetBackup for OpenStack dashboard provides cluster-wide management capability.

Note: For NetBackup Flex Appliance, you must manually copy the valid NetBackup CA-signed certificates at the `/etc/nbos/ssl/nbu_cacert.pem` before you configure NetBackup for OpenStack.

Details needed for the NetBackup for OpenStack Appliance

When you log in to an unconfigured NetBackup for OpenStack Appliance, the shown page is the configurator. The configurator requires some information about the NetBackup for OpenStack Appliance and OpenStack.

NetBackup for OpenStack Cluster information

The NetBackup for OpenStack Cluster needs to be integrated into an existing environment to be able to operate correctly. This block asks for information about the NetBackup for OpenStack Cluster operating details.

- Controller Nodes
 - This is the list of NetBackup for OpenStack virtual appliance IP addresses along with their host names.
 - Format: comma-separated list with pairs combined through "=". The first node in this list must be an active node.
 - Example:
172.20.4.151=nbos-104-1,172.20.4.152=nbos-104-2,172.20.4.153=nbos-104-3'

The NetBackup for OpenStack Cluster supports only 1-node and 3-node clusters.

- Virtual IP address
 - NetBackup for OpenStack cluster IP address, which is mandatory.
 - Format: IP/Subnet
 - Example: 172.20.4.150/24

Warning: The Virtual IP is mandatory even for single-node clusters and has to be different from any IP given at the Controller Nodes.

- Name Server
 - List of nameservers, primarily used to resolve OpenStack service endpoints.
 - Format: Comma-separated list
 - Example: 8.8.8.8,172.20.4.1
- Domain Search Order

- The domain the NetBackup for OpenStack Cluster will use.
- Format: Comma-separated list
- Example: nbos.io, nbos.demo
- NTP Servers
 - NTP servers the NetBackup for OpenStack Cluster will use
 - Format: Comma-separated list
 - Example: 0.pool.ntp.org,10.10.10.10
- Timezone
 - Timezone the NetBackup for OpenStack Cluster will use internally
 - Format: pre-populated list
 - Example: UTC

OpenStack Credentials information

The NetBackup for OpenStack appliance integrates with one RHV environment. This block asks for the information that is required to access and connect with the RHV Cluster.

- Keystone URL
 - The Keystone endpoint that is used to fetch authentication for configuration.
 - Format: URL
 - Example: `https://keystone.nbos.io:5000/v3`
- Endpoint Type
 - Defines which endpoint type is used to communicate with the OpenStack endpoints. NetBackup for OpenStack supports public endpoints and internal endpoints.
 - Format: Predefined list of radio buttons
 - Example: Public

When FQDNs are used for the Keystone endpoints it is necessary to configure at least one DNS server before the configuration.

Otherwise, the validation of the OpenStack Credentials fails.

- Domain ID
 - The domain of the provided user and the tenant
 - Format: ID

- Example: Default
- Administrator
 - Username of an account with the domain admin role
 - Format: String
 - Example: admin
- Password
 - Password for the previous provided user
 - Format: String
 - Example: Password

NetBackup for OpenStack requires domain admin role access. To provide a domain admin role to a user, the following command can be used:

```
openstack role add --domain <domain id> --user <username> admin
```

The NetBackup for OpenStack configurator verifies after every entry if it is possible to log in to OpenStack using the provided credentials.

This verification fails until all entries are set and correct.

When the verification is successful it is possible to choose the Admin tenant, the Region, and the Trustee role without any error message shown.

- Admin Tenant
 - The tenant to be used together with the provided user.
 - Format: A pre-populated list
 - Example: Admin
- Region
 - OpenStack Region the user and tenant are located in.
 - Format: a pre-populated list
 - Example: RegionOne
- Trustee Role
 - The OpenStack role is required to be able to use NetBackup for OpenStack functions.
 - Format: A pre-populated list
 - Example: _member_

Advanced settings

At the end of the configurator, you can activate the advanced settings.

Activating this option enables the configuration of the keystone endpoints that are used for NetBackup for OpenStack job manager and NetBackup for OpenStack datamover API.

Set up the NetBackup for OpenStack job manager and NetBackup for OpenStack datamover API

NetBackup for OpenStack generates keystone endpoints for two services. The NetBackup for OpenStack datamover API and the NetBackup for OpenStack job manager.

Modern OpenStack installation has the endpoint types split over multiple networks. The advanced settings for the nbosdmapi endpoints and nbosjm endpoints allow configuring NetBackup for OpenStack accordingly.

Used IP addresses are added as additional VIPs to the NetBackup for OpenStack cluster.

In the case of FQDN used for those endpoints the NetBackup for OpenStack configurator resolves the FQDN to learn the IPs that are then set as VIPs.

It is recommended to verify the nbosdmapi settings against the settings configured during installation of the NetBackup for OpenStack components.

If these endpoints do already exist in keystone the values are pre-filled and cannot be changed. In case of a change required, delete the old keystone endpoints first.

Providing a URL with https activates the TLS enabled configuration, which requires the upload of certificates and the connected private key.

See [“Configuring the secure communication for NBOSVM”](#) on page 59.

Configuring the secure communication for NBOSVM

You can use the secure communication for NBOSVM. If you want to use the secure communication, you must upload the certificate and its private key.

To configure the secure communication

- 1 Log on to the NetBackup for OpenStack configurator UI.
- 2 On the **Configuration Details** tab, click **Reconfigure** at the end of the page.
- 3 Click **Advanced Settings**.
- 4 In the **NetBackup for OpenStack URL(s)** field, change **HTTP** to **HTTPS** in the URL.

5 Click **Certificate** and **Private Key** to upload the files.

To generate the certificate and the key, go to the location `/etc/nbos/ssl` on any of the NBOSVM nodes and run the following command:

```
./gen-cer <NBOSVM VIP>
```

This command generates the certificate and key files with NBOSVM virtual IP as a file name.

For example, if the NBOSVM virtual IP is 10.10.20.111, you run the command `./gen-cert 10.10.20.111`

This command generates files such as `10.10.20.111.crt` and `10.10.20.111.key`.

Upload the `10.10.20.111.crt` and `10.10.20.111.key` files.

6 Click the drop-down next to the **NetBackup for OpenStack URL(s)** field.

In the **NetBackup for OpenStack Admin URL** and **NetBackup for OpenStack Internal URL** fields, change **HTTP** to **HTTPS**.

7 After NBOSVM configuration is successful, copy `/opt/stack/data/cert/nbosjm.cert` file from NBOSVM to each controller node at the following location.

- Kolla-openstack: `/etc/kolla/horizon`
- RHOSP: `/var/lib/config-data/puppet-generated/horizon`

8 Provide the following permissions to these files.

- Kolla-openstack:

```
chmod o+x /etc/kolla/horizon
chmod o+rx /etc/kolla/horizon/nbosjm.cert
```

- RHOSP:

```
chmod o+rx
/var/lib/config-data/puppet-generated/horizon/nbosjm.cert
```

9 Run the following command on the NBOSVM before you use the nbosjm CLI.

```
export OS_CACERT=/etc/nbosjm/ca-chain.pem
```

Set up an external database

NetBackup for OpenStack allows the use of an external MySQL or MariaDB database.

This database needs to be prepared by creating the empty nbosjm database, creating the nbosjm user and setting the right permissions. An example command to create this database would be:

```
create database nbosjm_auto;  
CREATE USER 'nbos'@'localhost' IDENTIFIED BY 'password';  
GRANT ALL PRIVILEGES ON nbosjm_auto.* TO 'nbos'@'10.10.10.67'  
IDENTIFIED BY 'password';
```

Provide the connection string to the NetBackup for OpenStack configurator.

```
mysql://nbos:password@10.10.10.67/nbosjm_auto?charset=utf8
```

This value can only be set upon an initial configuration of the NetBackup for OpenStack solution.

When the Cluster has been configured to use the internal database, then the connection string will not be shown in the next configuration attempt.

In the case of an external database, the connection string is shown but is not editable.

Define the NetBackup for OpenStack service user password

NetBackup for OpenStack is using a service user that is located in the OpenStack service project.

The password for this service user will be generated randomly or can be defined in the advanced settings.

Starting the configurator

Once all entries have been set and all validations are error-free the configurator can be started.

- Click Finish.
- Reconfirm in the pop-up that you want to start the configuration.
- Wait for the configurator to finish.

Some elements of the configurator take time. Even when it looks like the configurator is stuck, wait till the configurator finishes. If the configurator does not finish after 6 hours, contact Cohesity Support for help.

The configurator is using Ansible and a few NetBackup for OpenStack internal API calls. After each configuration block or after the configurator finished it is possible to visit the Ansible output.

At the end of a successful configuration the configurator will redirect the NBOSVM dashboard to virtual IP.

Resource throttling in NetBackup for OpenStack

Resource throttling is done to manage performance, prevent system overload, and ensure the fair NetBackup for OpenStack resource allocation among projects. Throttling prevents excessive consumption of resources and helps maintain stability. Throttling also prevents system failures and the issues that occur due to degraded performance.

Table 2-6 Resource throttling options in NetBackup for OpenStack

Option	Description
<code>MAX_BFS_JOBS_PER_NBOS</code>	<p>Resource throttling at the NetBackup side to concurrently run the number of backups from snapshot jobs.</p> <p>Configure the value in the <code>/usr/opensv/netbackup/bp.conf</code> file on the NetBackup primary server.</p> <p>Default value: <code>MAX_BFS_JOBS_PER_NBOS = 3</code></p>
<code>max_snapshot_jobs_per_project</code>	<p>Resource throttling on NetBackup for OpenStack virtual machine to concurrently run the number of snapshot jobs per project.</p> <p>Configure the value in the <code>/etc/nbosjm/nbosjm.conf</code> file on each nbosvm node.</p> <p>Restart the <code>nbosjm-scheduler</code> service on one of the nbosvm nodes where this service is running. On a three-node cluster, run the following command to know on which nbosvm node the <code>nbosjm-scheduler</code> service is running: <code>pcs status</code></p> <p>Default value: <code>max_snapshot_jobs_per_project = 2</code></p>

Table 2-6 Resource throttling options in NetBackup for OpenStack
(continued)

Option	Description
<code>max_snapshot_expiry_jobs_per_project</code>	<p>Resource throttling on NetBackup for OpenStack virtual machine to concurrently expire the number of snapshot jobs per project.</p> <p>Configure the value in the <code>/etc/nbosjm/nbosjm.conf</code> file on each nbosvm node.</p> <p>Restart the <code>nbosjm-scheduler</code> service on one of the nbosvm nodes where this service is running. On a three-node cluster, run the following command to know on which nbosvm node the <code>nbosjm-scheduler</code> service is running: <code>pcs status</code></p> <p>Default value: <code>max_snapshot_expiry_jobs_per_project</code> <code>= 2</code></p>
<code>max_uploads_pending</code>	<p>Resource throttling on NetBackup for OpenStack datamover to concurrently run data movement operation per compute node.</p> <p>Configure the value in <code>nbosdm.conf</code> file on the compute node.</p> <p>Restart the nbosdm container running on that compute node.</p> <p>Default value: <code>max_uploads_pending</code> = <code>5</code></p>

Post Installation Health-Check

After the installation and configuration of NetBackup for OpenStack did succeed the following steps can be done to verify that the NetBackup for OpenStack installation is healthy.

Verify the NetBackup for OpenStack Appliance services are up

NetBackup for OpenStack uses three main services on the NetBackup for OpenStack Appliance:

- nbosjm-api
- nbosjm-scheduler
- nbosjm-policies

Those can be verified to be up and running using the `systemctl status` command.

```
systemctl status nbosjm-api
#####
● nbosjm-api.service - nbosjm api service
   Loaded: loaded (/etc/systemd/system/nbosjm-api.service; disabled;
           vendor preset: disabled)
   Drop-In: /run/systemd/system/nbosjm-api.service.d
            └─50-pacemaker.conf
   Active: active (running) since Wed 2020-04-22 09:17:05 UTC; 1 day 2h ago
   Main PID: 21265 (python)
     Tasks: 1
    CGroup: /system.slice/nbosjm-api.service
            └─21265 /home/rhv/myansible/bin/python /usr/bin/nbosjm-api
              --config-file=/etc/nbosjm/nbosjm.conf

systemctl status nbosjm-scheduler
#####
● nbosjm-scheduler.service - nbosjm scheduler service
   Loaded: loaded (/etc/systemd/system/nbosjm-scheduler.service; disabled;
           vendor preset: disabled)
   Drop-In: /run/systemd/system/nbosjm-scheduler.service.d
            └─50-pacemaker.conf
   Active: active (running) since Wed 2020-04-22 09:17:17 UTC; 1 day 2h ago
   Main PID: 21512 (python)
     Tasks: 1
    CGroup: /system.slice/nbosjm-scheduler.service
            └─21512 /home/rhv/myansible/bin/python /usr/bin/nbosjm-scheduler
              --config-file=/etc/nbosjm/nbosjm.conf

systemctl status nbosjm-policies
#####
● nbosjm-policies.service - nbosjm policies service
   Loaded: loaded (/etc/systemd/system/nbosjm-policies.service; enabled;
           vendor preset: disabled)
   Active: active (running) since Wed 2020-04-22 09:15:43 UTC; 1 day 2h ago
   Main PID: 20079 (python)
```



```

Tasks: 33
CGroup: /system.slice/nbosjm-policies.service
└─20079 /home/rhv/myansible/bin/python
/usr/bin/nbosjm-policies
--config-file=/etc/nbosjm/nbosjm.conf
└─20180 /home/rhv/myansible/bin/python
/usr/bin/nbosjm-policies
--config-file=/etc/nbosjm/nbosjm.conf
[...]
└─20181 /home/rhv/myansible/bin/python
/usr/bin/nbosjm-policies
--config-file=/etc/nbosjm/nbosjm.conf
└─20233 /home/rhv/myansible/bin/python
/usr/bin/nbosjm-policies
--config-file=/etc/nbosjm/nbosjm.conf
└─20236 /home/rhv/myansible/bin/python
/usr/bin/nbosjm-policies
--config-file=/etc/nbosjm/nbosjm.conf
└─20237 /home/rhv/myansible/bin/python
/usr/bin/nbosjm-policies
--config-file=/etc/nbosjm/nbosjm.conf

```

Check the NetBackup for OpenStack pacemaker and NGINX cluster

The second component to check the NetBackup for OpenStack Appliance's health is the NGINX and pacemaker cluster.

```

pcs status
#####
Cluster name: NetBackup for OpenStack

WARNINGS:
Corosync and pacemaker node names do not match (IPs used in setup?)
Stack: corosync
Current DC: om_nbosvm (version 1.1.19-8.el7_6.1-c3c624ea3d) -
chapterition with quorum
Last updated: Wed Dec 5 12:25:02 2018
Last change: Wed Dec 5 09:20:08 2018 by root via cibadmin on om_nbosvm
1 node configured
4 resources configured

Online: [ om_nbosvm ]

```

```
Full list of resources:
virtual_ip (ocf::'heartbeat:IPaddr2): Started om_nbosvm
nbosjm-api (systemd:nbosjm-api): Started om_nbosvm
nbosjm-scheduler (systemd:nbosjm-scheduler): Started om_nbosvm
Clone Set: lb_nginx-clone [lb_nginx]
Started: [ om_nbosvm ]
Daemon Status:
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
```

Verify API connectivity of the NetBackup for OpenStack Appliance

Checking the availability of the NetBackup for OpenStack API on the chosen endpoints is recommended.

The following example curl command lists the available policy types and verifies that the connection is available and working:

```
curl http://10.10.2.34:8780/v1/8e16700ae3614da4ba80a4e57d60cdb9/
policy_types/detail -X GET -H "X-Auth-Project-Id: admin"
-H "User-Agent: python-nbosjmclient" -H "Accept:
application/json" -H "X-Auth-Token:
gAAAAABe40NVFETJeePpk1F9QGGh1LiGnHJVLlgZx9t0HRRK9rC5vq
KZJRkpAcWloPH6Q9K9peuHiQrBHEs1-g75Na4xOEESR0LmQJUZF6n3
7fLfDL_D-hlnjHJZ68iNisIPlfkm9FGSyoyt6IqjO9E7_YVRCTCqNLJ
67ZkqHuJh1CXwShvjvfw
```

See the API guide for more commands and to know how to generate the X-Auth-Token.

Verify that the nbosdm services are up and running

The nbosdm service is the datamover that got installed on all compute nodes. Check its status after the installation.

```
[root@upstreamcompute1 ~]# systemctl status tripleo-nbosdm.service
• tripleo_nbosdm.service - nbosdm container
   Loaded: loaded (/etc/systemd/system/tripleo_nbosdm.service; enabled;
          vendor preset: disabled)
   Active: active (running) since Wed 2020-06-10 10:07:28 EDT; 1 day 19h ago
   Main PID: 10384 (python)
   Tasks: 21
```

```

CGroup: /system.slice/tripleo_nbosdm.service
└─10384 /usr/bin/python /usr/bin/nbosdm --config-file=/etc...

Jun 12 03:15:33 upstreamcompute1 python[10384]: libvirt: QEMU Driver
error :...d
Jun 12 03:15:33 upstreamcompute1 python[10384]: libvirt: QEMU Driver
error :...d
Jun 12 03:16:11 upstreamcompute1 python[10384]: libvirt: QEMU Driver
error :...d
Jun 12 03:16:31 upstreamcompute1 sudo[13977]:      nova : TTY=unknown ;
PWD=/...n
Jun 12 03:16:33 upstreamcompute1 sudo[14004]:      nova : TTY=unknown ;
PWD=/ ...
Jun 12 05:15:33 upstreamcompute1 python[10384]: libvirt: QEMU Driver
error :...d
Jun 12 05:15:33 upstreamcompute1 python[10384]: libvirt: QEMU Driver
error :...d
Jun 12 05:16:11 upstreamcompute1 python[10384]: libvirt: QEMU Driver
error :...d
Jun 12 05:16:29 upstreamcompute1 sudo[23356]:      nova : TTY=unknown ;
PWD=/...n
Jun 12 05:16:32 upstreamcompute1 sudo[23422]:      nova : TTY=unknown ;
PWD=/ ...
Hint: Some lines were ellipsized, use -l to show in full.

```

Uninstalling NetBackup for OpenStack

The steps to uninstall NetBackup for OpenStack vary depending on the OpenStack distribution it is installed in. However, the following high-level steps are the same for all distributions.

1. Uninstall the Horizon plug-in or the NetBackup OpenStack Horizon container.
2. Uninstall the nbosdmapi container.
3. Uninstall the nbosdm.
4. Delete the NetBackup for OpenStack Cluster.

Uninstalling from RHOSP

Perform the following steps to uninstall NetBackup for OpenStack from RHOSP:

Clean the NetBackup for OpenStack
datamover API service.

See [“Clean NetBackup for OpenStack
datamover API service”](#) on page 68.

Clean the NetBackup for OpenStack datamover service.	See “Clean NetBackup for OpenStack datamover service” on page 69.
Clean the NetBackup for OpenStack haproxy resources.	See “Clean NetBackup for OpenStack haproxy resources” on page 70.
Clean the NetBackup for OpenStack Keystone resources.	See “Clean NetBackup for OpenStack Keystone resources” on page 71.
Clean the NetBackup for OpenStack database resources.	See “Clean NetBackup for OpenStack database resources” on page 71.
Revert the overcloud deploy command.	See “Revert overcloud deploy command” on page 72.
Revert back to the original RHOSP Horizon container.	See “Revert back to original RHOSP Horizon container” on page 73.
Destroy the NetBackup for OpenStack virtual machine cluster.	See “Destroy the NetBackup for OpenStack virtual machine cluster” on page 73.

Clean NetBackup for OpenStack datamover API service

The following steps need to be run on all nodes, which have the NetBackup for OpenStack datamover API service running. Those nodes can be identified by verifying the `roles_data.yaml` for the role that contains the entry

```
OS::TripleO::Services::nbosdmapi.
```

Once the role that runs the NetBackup for OpenStack datamover API service has been identified, the following commands will clean the nodes from the service.

Warning: Run all commands as root or user with sudo permissions.

Stop `nbosdmapi` container.

```
# For RHOSP17.1 onwards
systemctl disable tripleo_nbosdmapi.service
systemctl stop tripleo_nbosdmapi.service
podman stop nbosdmapi
```

Remove `nbosdmapi` container.

```
# For RHOSP17.1 onwards
podman rm nbosdmapi
```

```
podman rm nbosdmap_i_init_log
podman rm nbosdmap_i_db_sync
```

Clean NetBackup for OpenStack datamover API service conf directory.

```
rm -rf /var/lib/config-data/puppet-generated/nbosdmap_i
rm /var/lib/config-data/puppet-generated/nbosdmap_i.md5sum
```

Clean NetBackup for OpenStack datamover API service log directory.

```
rm -rf /var/log/containers/nbosdmap_i/
```

Clean NetBackup for OpenStack datamover service

The following steps need to be run on all nodes, which have the NetBackup for OpenStack datamover service running. Those nodes can be identified by checking the `roles_data.yaml` for the role that contains the entry

```
OS::TripleO::Services::nbosdm.
```

Once the role that runs the NetBackup for OpenStack datamover API service has been identified, the following commands will clean the nodes from the service.

Warning: Run all commands as root or user with sudo permissions.

Stop `nbosdm` container.

```
# For RHOSP17.1 onwards
systemctl disable tripleo_nbosdm.service
systemctl stop tripleo_nbosdm.service
podman stop nbosdm
```

Remove `nbosdm` container.

```
# For RHOSP17.1 onwards
podman rm nbosdm
```

Unmount NetBackup for OpenStack Backup Target on compute host.

```
## Following steps applicable for all supported RHOSP releases.
```

```
# Check NetBackup for OpenStack backup target mount point
```

```
mount | grep NetBackup

# Unmount it
-- If it's NFS (COPY UUID_DIR from your compute host using above command)
umount /var/lib/nova/NetBackupOpenStack-mounts/<UUID_DIR>

-- If it's S3
umount /var/lib/nova/NetBackupOpenStack-mounts

# Verify that it's unmounted
mount | grep NetBackup

df -h | grep NetBackup

# Remove mount point directory after verifying that backup target unmounted
successfully.
# Otherwise actual data from backup target may get cleaned.

rm -rf /var/lib/nova/NetBackupOpenStack-mounts
```

Clean NetBackup for OpenStack datamover service conf directory.

```
rm -rf /var/lib/config-data/puppet-generated/nbosdm/
rm /var/lib/config-data/puppet-generated/nbosdm.md5sum
```

Clean log directory of NetBackup for OpenStack datamover service.

```
rm -rf /var/log/containers/nbosdm/
```

Clean NetBackup for OpenStack haproxy resources

The following steps need to be run on all nodes, which have the haproxy service running. Those nodes can be identified by verifying the `roles_data.yaml` for the role that contains the entry `OS::TripleO::Services::Haproxy`.

Once the role that runs the NetBackup for OpenStack datamover API service has been identified, the following commands will clean the nodes from all NetBackup for OpenStack resources..

Warning: Run all commands as root or user with sudo permissions.

Edit the following file on the HAProxy nodes and remove all NetBackup for OpenStack entries.

`/var/lib/config-data/puppet-generated/haproxy/etc/haproxy/haproxy.cfg`

An example of these entries:

```
listen nbosdmapi
  bind 172.25.3.60:13784 transparent ssl crt /etc/pki/tls/private/
  overcloud_endpoint.pem
  bind 172.25.3.60:8784 transparent
  http-request set-header X-Forwarded-Proto https if { ssl_fc }
  http-request set-header X-Forwarded-Proto http if !{ ssl_fc }
  http-request set-header X-Forwarded-Port %[dst_port]
  option httpchk
  option httplog
  server overcloud-controller-0.internalapi.localdomain 172.25.3.59:8784
  check fall 5 inter 2000 rise 2
```

Restart the haproxy container once all edits have been done.

```
# For RHOSP17.1 onwards
podman restart haproxy-bundle-podman-0
```

Clean NetBackup for OpenStack Keystone resources

NetBackup for OpenStack registers services and users in Keystone. Those need to be unregistered and deleted.

```
openstack service delete nbosdmapi
openstack user delete nbosdmapi
```

Clean NetBackup for OpenStack database resources

NetBackup for OpenStack creates a database for the nbosdmapi service. This database needs to be cleaned.

Login into the database cluster.

```
## On RHOSP
podman exec -it galera-bundle-podman-0 mysql -u root
```

Run the following SQL statements to clean the database.

```
## Clean database
DROP DATABASE nbosdmapi;

## Clean nbosdmapi user
MariaDB [mysql]> select user, host from mysql.user where user='nbosdmapi';
+-----+-----+
| user      | host      |
+-----+-----+
| nbosdmapi | 172.25.2.10 |
| nbosdmapi | 172.25.2.8  |
+-----+-----+
2 rows in set (0.00 sec)

=> Delete those user accounts
MariaDB [mysql]> DROP USER nbosdmapi@172.25.2.10;
Query OK, 0 rows affected (0.82 sec)

MariaDB [mysql]> DROP USER nbosdmapi@172.25.2.8;
Query OK, 0 rows affected (0.05 sec)

=> Verify that nbosdmapi user got cleaned
MariaDB [mysql]> select user, host from mysql.user where user='nbosdmapi';
Empty set (0.00 sec)
```

Revert overcloud deploy command

Remove the following entries from `roles_data.yaml` used in the overcloud deploy command.

- `OS::TripleO::Services::nbosdmapi`
- `OS::TripleO::Services::nbosdm`

In case the overcloud deploy command used before the deployment of NetBackup for OpenStack is still available, it can directly be used.

Follow these steps to clean the overcloud deploy command from all NetBackup for OpenStack entries.

1. Remove `nbos_env.yaml` entry.
2. Remove NetBackup OpenStack endpoint map file Replace with original map file if existing.

Revert back to original RHOSP Horizon container

Run the cleaned overcloud deploy command.

Destroy the NetBackup for OpenStack virtual machine cluster

List all virtual machines running on the KVM node

```
virsh list
```

Destroy the NetBackup for OpenStack virtual machines

```
virsh destroy <NetBackup for OpenStack virtual machine Name or ID>
```

Undefine the NetBackup for OpenStack virtual machines

```
virsh undefine <NetBackup for OpenStack virtual machine name>
```

Delete the NetBackup for OpenStack virtual machine disk from KVM Host storage

Uninstalling from Ansible OpenStack

Perform the following tasks to uninstall NetBackup for OpenStack from Ansible OpenStack:

Uninstall NetBackup for OpenStack Services See [“Uninstall NetBackup for OpenStack Services”](#) on page 74.

Destroy NetBackup for OpenStack datamover API container See [“Destroy NetBackup for OpenStack datamover API container”](#) on page 74.

Clean openstack_user_config.yml See [“Clean openstack_user_config.yml”](#) on page 74.

Remove NetBackup for OpenStack haproxy settings in user_variables.yml See [“Remove NetBackup for OpenStack haproxy settings in user_variables.yml”](#) on page 75.

Remove NetBackup for OpenStack datamover API inventory file See [“Remove NetBackup for OpenStack datamover API inventory file”](#) on page 75.

Remove NetBackup for OpenStack datamover API service endpoints See [“Remove NetBackup for OpenStack datamover API service endpoints”](#) on page 75.

Delete NetBackup for OpenStack datamover API database and user	See “Delete NetBackup for OpenStack datamover API database and user” on page 76.
Remove nbosdmapi rabbitmq user from rabbitmq container	See “Remove nbosdmapi rabbitmq user from rabbitmq container” on page 76.
Clean haproxy	See “Clean haproxy” on page 76.
Remove certificates from Compute nodes	See “Remove certificates from Compute nodes” on page 77.
Destroy the NetBackup for OpenStack virtual machine cluster	See “Destroy the NetBackup for OpenStack virtual machine cluster” on page 77.

Uninstall NetBackup for OpenStack Services

The NetBackup for OpenStack Ansible OpenStack playbook can be run to uninstall the NetBackup for OpenStack services.

```
cd /opt/openstack-ansible/playbooks
openstack-ansible os-nbos-install.yml --tags "nbos-all-uninstall"
```

Destroy NetBackup for OpenStack datamover API container

To cleanly remove the NetBackup for OpenStack datamover API container run the following Ansible playbook.

```
cd /opt/openstack-ansible/playbooks
openstack-ansible lxc-containers-destroy.yml --limit "DMPAI CONTAINER_NAME"
```

Clean openstack_user_config.yml

Remove the `nbosdmapi_hosts` and `nbos_compute_hosts` entries from `/etc/openstack_deploy/openstack_user_config.yml`

```
#nbosdmapi
nbos-nbosdmapi_hosts:
  infra-1:
    ip: 172.26.0.3
  infra-2:
    ip: 172.26.0.4
```

```
#nbos-datamover
nbos_compute_hosts:
  infra-1:
    ip: 172.26.0.7
  infra-2:
    ip: 172.26.0.8
```

Remove NetBackup for OpenStack haproxy settings in user_variables.yml

Remove NetBackup for OpenStack datamover API settings from
/etc/openstack_deploy/user_variables.yml

```
# Datamover haproxy setting
haproxy_extra_services:
  - service:
      haproxy_service_name: nbosdm_service
      haproxy_backend_nodes: "{{ groups['nbosdmapi_all'] | default([]) }}"
      haproxy_ssl: "{{ haproxy_ssl }}"
      haproxy_port: 8784
      haproxy_balance_type: http
      haproxy_balance_alg: roundrobin
      haproxy_timeout_client: 10m
      haproxy_timeout_server: 10m
      haproxy_backend_options:
        - "httpchk GET / HTTP/1.0\\r\\nUser-agent:\\\\ osa-haproxy-healthcheck"
```

Remove NetBackup for OpenStack datamover API inventory file

```
rm /opt/openstack-ansible/inventory/env.d/nbos-nbosdmapi.yml
```

Remove NetBackup for OpenStack datamover API service endpoints

```
source cloudadmin.rc
openstack endpoint delete "internal datamover service endpoint_id"
openstack endpoint delete "public datamover service endpoint_id"
openstack endpoint delete "admin datamover service endpoint_id"
```

Delete NetBackup for OpenStack datamover API database and user

- Go inside galera container.
- Login as root user in mysql database engine.
- Drop nbosdmapi database.
- Drop nbosdmapi user

```
lxc-attach -n "GALERA CONTAINER NAME"
mysql -u root -p "root password"
DROP DATABASE nbosdmapi;
DROP USER nbosdmapi;
```

Remove nbosdmapi rabbitmq user from rabbitmq container

- Go inside rabbitmq container.
- Delete nbosdmapi user.
- Delete nbosdmapi vhost.

```
lxc-attach -n "RABBITMQ CONTAINER NAME"
rabbitmqctl delete_user nbosdmapi
rabbitmqctl delete_vhost /nbosdmapi
```

Clean haproxy

Remove `/etc/haproxy/conf.d/nbosdm_service` file.

```
rm /etc/haproxy/conf.d/nbosdm_service
```

Remove HAProxy configuration entry from `/etc/haproxy/haproxy.cfg` file.

```
frontend nbosdm_service-front-1
    bind hostname:8784 ssl crt /etc/ssl/private/
    haproxy.pem ciphers ECDH+AESGCM:DH+AESGCM:ECDH
    +AES256:DH+AES256:ECDH+AES128:DH+AES:RSA+AESGCM
    :RSA+AES:!aNULL:!MD5:!DSS
    option httplog
    option forwardfor except 127.0.0.0/8
    reqadd X-Forwarded-Proto:\ https
```

```

mode http
default_backend nbosdm_service-back

frontend nbosdm_service-front-2
bind 172.26.1.2:8784
option httplog
option forwardfor except 127.0.0.0/8
mode http
default_backend nbosdm_service-back

backend nbosdm_service-back
mode http
balance leastconn
stick store-request src
stick-table type ip size 256k expire 30m
option forwardfor
option httplog
option httpchk GET / HTTP/1.0\r\nUser-agent:\ osa-haproxy-healthcheck

server controller_nbosdmapi_container-bf17d5b3 172.26.1.75:8784
check port 8784 inter 12000 rise 1 fall 1

```

Restart the HAproxy service.

```
systemctl restart haproxy
```

Remove certificates from Compute nodes

```
rm -rf /opt/config-certs/rabbitmq
rm -rf /opt/config-certs/s3
```

Destroy the NetBackup for OpenStack virtual machine cluster

List all virtual machines running on the KVM node

```
virsh list
```

Destroy the NetBackup for OpenStack virtual machines

```
virsh destroy <NetBackup for OpenStack virtual machine Name or ID>
```

Undefine the NetBackup for OpenStack virtual machines

```
virsh undefine <NetBackup for OpenStack virtual machine name>
```

Delete the NetBackup for OpenStack virtual machine disk from KVM Host storage

Uninstalling from Kolla Openstack

Cleaning NetBackupOpenStack_datamover_api container

The container needs to be cleaned on all nodes where the NetBackupOpenStack_datamover_api container is running. The Kolla Openstack inventory file helps to identify the nodes with the service.

Following steps need to be done to clean the NetBackupOpenStack_datamover_api container:

To clean NetBackupOpenStack_datamover_api container

- 1 Stop the NetBackupOpenStack_datamover_api container.

```
docker stop NetBackupOpenStack_datamover_api
```

- 2 Remove the NetBackupOpenStack_datamover_api container.

```
docker rm NetBackupOpenStack_datamover_api
```

- 3 Clean /etc/kolla/nbosdmapapi directory.

```
rm -rf /etc/kolla/nbosdmapapi
```

- 4 Clean log directory of NetBackupOpenStack_datamover_api container.

```
rm -rf /var/log/kolla/nbosdmapapi/
```

Cleaning NetBackupOpenStack_datamover container

The container needs to be cleaned on all nodes where the NetBackupOpenStack_datamover container is running. The Kolla Openstack inventory file helps to identify the nodes with the service.

To clean NetBackupOpenStack_datamover container

- 1 Stop the NetBackupOpenStack_datamover container.

```
docker stop NetBackupOpenStack_datamover
```

- 2 Remove the NetBackupOpenStack_datamover container.

```
docker rm NetBackupOpenStack_datamover
```

- 3 Clean /etc/kolla/nbosdm directory.

```
rm -rf /etc/kolla/nbosdm
```

- 4 Clean log directory of NetBackupOpenStack_datamover container.

```
rm -rf /var/log/kolla/nbosdm/
```

Cleaning haproxy of NetBackupOpenStack datamover API

The NetBackupOpenStack Datamover API entries need to be cleaned on all `haproxy` nodes. The Kolla Openstack inventory file helps to identify the nodes with the service.

To clean haproxy of NetBackupOpenStack datamover API

- 1

```
rm /etc/kolla/haproxy/services.d/nbosdmapl.cfg
```

- 2

```
docker restart haproxy
```

Cleaning Kolla Ansible deployment procedure

Delete all NetBackup for OpenStack related entries from:

- `/path/to/venv/share/kolla-ansible/ansible/roles/` There is a role NetBackup for OpenStack.
- `/etc/kolla/globals.yml` NetBackup for OpenStack entries are appended at the end of the file.
- `/etc/kolla/passwords.yml` NetBackup for OpenStack entries had been appended at the end of the file.
- `/path/to/venv/share/kolla-ansible/ansible/site.yml` NetBackup for OpenStack entries had been appended at the end of the file.
- `/root/multinode` NetBackup for OpenStack entries are appended at the end of this example inventory file.

Reverting to original Horizon container

Run deploy command to replace the NetBackup for OpenStack Horizon container with original Kolla Ansible Horizon container.

```
kolla-ansible -i multinode deploy
```

Cleaning Keystone resources

NetBackup for OpenStack created a nbosdmapi service with nbosdmapi user. Run the following commands to clean Keystone resources.

```
openstack service delete nbosdmapi
```

```
openstack user delete nbosdmapi
```

Cleaning NetBackup for OpenStack database resources

NetBackup for OpenStack datamover API service has its own database in the OpenStack database.

To clean NetBackup for OpenStack database resources

- 1 Login to Openstack database as root user or user with similar privileges.

```
mysql -u root -p
```

- 2 Delete nbosdmapi database and user.

```
DROP DATABASE nbosdmapi;
```

```
DROP USER nbosdmapi;
```

Destroy the NetBackup for OpenStack virtual machine cluster

To destroy the NetBackup for OpenStack virtual machine cluster

- 1 List all virtual machines running on the KVM node

```
virsh list
```

- 2 Destroy the NetBackup for OpenStack virtual machines

```
virsh destroy <NetBackup for OpenStack virtual machine Name or ID>
```


3 Undefine the NetBackup for OpenStack virtual machines

```
virsh undefine <NetBackup for OpenStack virtual machine name>
```

4 Delete the NetBackup for OpenStack virtual machine disk from KVM Host storage.

5 Deregister the NetBackup for OpenStack.

- Generate a token using the following API:

```
POST https://<primary-server>:1556/netbackup/login
```

Provide username and password in body as :

```
{
  "userName": "username",
  "password": "password"
}
```

- Use the token in the following de-register API as a bearer token:

```
DELETE
```

```
https://<primary-server>/netbackup/config/servers/nbosvm-servers/<NBOSVM
VIP
```

Install nbosjm CLI client

The `nbosjm` CLI client is provided as RPM and Debian packages. Installing the `nbosjm` client automatically installs all the required OpenStack clients.

The installation of the `nbosjm` client integrates the client into the global OpenStack python client, if available.

The required connection strings and package names can be found on the NetBackup for OpenStack dashboard under the Downloads tab.

The `nbosjm` client is supported only on Python 3.

To install the `nbosjm` CLI client

- ◆ ■ Run the following command on RPM-based operating systems:

```
yum install nbosjmclient-py3-el8-9.0.999-9.0.noarch.rpm
```
- Run the following command on Debian-based operating systems:

```
apt-get install nbosjmclient-py3_9.0.999_all.deb
```

About log rotation in NetBackup for OpenStack

Use log rotation to ease administration of the systems that generate large numbers of log files. It allows automatic rotation, compression, removal, and mailing of log files. You can handle each log file daily, weekly, monthly, or when it grows too large.

logrotate is a Linux utility, which runs as a scheduled cron job. It reads the information from the configuration files. You can configure log rotation by updating these configuration files.

On the RHOSP platform, after configuration changes for log rotation, you must update the stack for the changes to take effect.

Empty VxMS log files are cleaned up automatically after 8 days.

[Table 2-7](#) describes the default options that are used to configure log rotation on Kolla and Ansible.

Table 2-7 Log rotation default options for Kolla and Ansible

Component	Log rotation default options
NBOSJM	<p>Configuration file: /etc/logrotate.d/nbosjm</p> <pre> /var/log/nbosjm/*.log { missingok notifempty copytruncate size 25M rotate 30 compress dateformat -%Y%m%d-%H } /var/NetBackupOpenStack-mounts/*/*.log { su root nova missingok compress delaycompress notifempty copytruncate size 25M rotate 30 dateformat -%Y%m%d-%H } /var/NetBackupOpenStack-mounts/*/*.log { su root nova missingok compress delaycompress notifempty copytruncate size 25M rotate 30 dateformat -%Y%m%d-%H } </pre>

Table 2-7 Log rotation default options for Kolla and Ansible (*continued*)

Component	Log rotation default options
NBOSDMAPI	<p>Configuration file: <code>/etc/logrotate.d/nbosdmap</code></p> <pre> /var/log/kolla/nbosdmap/nbosdmap.log { daily missingok notifempty copytruncate size=25M rotate 30 maxage 30 compress dateformat -%Y%m%d-%H } </pre>
VxMS and NBOSDM	<p>Configuration file: <code>/etc/logrotate.d/nbosdm</code></p> <pre> /var/log/kolla/nbosdm/nbosdm.log { daily missingok notifempty copytruncate size=25M rotate 30 maxage 30 compress dateformat -%Y%m%d-%H } /usr/openv/netbackup/logs/vxms/*.log { daily missingok notifempty copytruncate size=1M rotate 5 postrotate find -daystart -mtime +30 -delete find -daystart -mtime +7 -size 0 -delete endscript compress dateformat -%Y%m%d-%H } </pre>

Table 2-8 describes the default options that are used to configure log rotation on RHOSP.

Table 2-8 Log rotation default options for RHOSP

Component	Log rotation default options
NBOSJM	<p>Configuration file: /etc/logrotate.d/nbosjm</p> <pre> /var/log/nbosjm/*.log { missingok notifempty copytruncate size 25M rotate 30 compress dateformat -%Y%m%d-%H } /var/NetBackupOpenStack-mounts/*/*/*.log { su root nova missingok compress delaycompress notifempty copytruncate size 25M rotate 30 dateformat -%Y%m%d-%H } /var/NetBackupOpenStack-mounts/*/*/*.log { su root nova missingok compress delaycompress notifempty copytruncate size 25M rotate 30 dateformat -%Y%m%d-%H } </pre>
NBOSDM and NBOSDMAPI	<p>Refer to the following file on the director node:</p> <pre> /home/stack/openstack-tripleo-heat-templates/deployment/ logrotate/logrotate-crond-container-puppet.yaml </pre>

Table 2-8 Log rotation default options for RHOSP (*continued*)

Component	Log rotation default options
VxMS	<pre>Configuration file: /etc/logrotate.d/nbosdm /etc/logrotate.d/vxms /var/log/vxms/*.log { daily missingok notifempty copytruncate size=1M rotate 5 postrotate find -daystart -mtime +30 -delete find -daystart -mtime +7 -size 0 -delete endscrip compress dateformat -%Y%m%d-%H }</pre>

Upgrading NetBackup for OpenStack

You can upgrade NetBackup for OpenStack from the previous release.

Prerequisites:

- Before you delete the existing NetBackup for OpenStack virtual machine, take the backup of the configuration file `/etc/nbosjm/nbosjm.conf`. You can use this file to apply the previous configuration changes to the new NetBackup for OpenStack virtual machine.
- All the snapshot jobs must be complete on the existing NetBackup for OpenStack virtual machine.
- All the backup jobs must be complete on the existing NetBackup for OpenStack virtual machine. Mark the incomplete SLP jobs as complete on the NetBackup primary server.
 - List the incomplete SLPs:
`nbstlutil stlilist -image_incomplete`
 - Cancel the incomplete SLPs:
`nbstlutil cancel -backupid`

- Delete the orphaned snapshots.
See [“Deleting the orphaned snapshots”](#) on page 88.

The upgrade does not preserve the following information:

- Snapshot only protections
- Email settings
- Audit logs
- Restored VMs data

The upgrade does not preserve the information from the following NetBackup for OpenStack database tables:

- `atomdetails`
- `auditlogs`
- `flowdetails`
- `logbooks`
- `restore_metadata`
- `restored_vm_meta`
- `restored_vm_res`
- `restored_vms restores`

To upgrade NetBackup for OpenStack

- 1 Download the latest NetBackup for OpenStack packages from the download center.
- 2 Spin up the new NetBackup for OpenStack virtual machine.
See [“Spinning up the NetBackup for OpenStack virtual machine”](#) on page 23.
- 3 Run the following command to import the protections on any one of the NetBackup for OpenStack virtual machines.

```
nbosjm protection-import
```

The protection import job time-out is 10 minutes by default. To configure the time-out specify `protection_import_job_timeout_in_mins=<minutes>` in the `/etc/nbosjm/nbosjm.conf` file.

- 4 Run the following command to see the status of the import job.

```
nbosjm get-protection-import-status
```

- 5 Start the global job scheduler.

```
nbosjm enable-global-job-scheduler
```

Deleting the orphaned snapshots

The recovery points that consist only snapshot copies are not imported during the upgrade process and these orphaned snapshots stay on the compute storage.

You can delete these snapshots using the `delete_snapshots.py` script.

To delete the orphaned snapshots

- 1 Copy `nbos-cfg-scripts` package to the old NBOSVM and extract the files.

```
cp <imagelocation>/nbos-cfg-scripts.tar.gz /home/stack  
tar -xvf /home/stack/nbos-cfg-scripts.tar
```

- 2 Go to the following directory.

```
cd /home/stack/nbos-cfg-scripts/
```

- 3 Activate the virtual environment of the nbosjm.

```
source /home/stack/myansible/bin/activate
```

- 4 Run the `delete_snapshots.py` script to delete the recovery points that consist only the snapshot copies.

```
python3 delete_snapshots.py delete
```

- 5 Run the following command to get the status of the snapshot deletion.

```
python3 delete_snapshots.py get_delete_status
```


Configuring NetBackup OpenStack Appliance

This chapter includes the following topics:

- [Reconfigure the NetBackup for OpenStack cluster](#)
- [Configuring the NetBackup primary server details](#)
- [Change NetBackup for OpenStack dashboard password](#)
- [Reset NetBackup for OpenStack dashboard password](#)
- [Downloading the NetBackup for OpenStack logs](#)
- [Updating the API key](#)
- [Uploading the API certificate](#)

Reconfigure the NetBackup for OpenStack cluster

The NetBackup for OpenStack appliance can be reconfigured at any time to adjust the NetBackup for OpenStack cluster to any changes in the OpenStack environment or the general backup solution.

To reconfigure the NetBackup for OpenStack cluster go to the "Configure". The configure page shows the current configuration of the NBOSVM cluster.

The configuration page also gives access to the Ansible playbooks of the last successful configuration.

To start the reconfiguration of the NetBackup for OpenStack cluster click **Reconfigure** at the end of the table.

Follow the Configuring NetBackup for OpenStack guide afterwards.

Once the NetBackup for OpenStack configurator has started, it needs to run through successfully to continue to use NetBackup for OpenStack.

The cluster does not roll back to its last working state in case of any errors.

Configuring the NetBackup primary server details

You must configure the primary server details on the NetBackup for OpenStack virtual machine. This configuration on the NetBackup for OpenStack configurator UI is required for the communication for license checks, capacity reporting, and certificate deployment.

To configure the primary server details

- 1 Log on to the NetBackup for OpenStack configurator UI.
- 2 Enter the primary server host name.
- 3 Enter the Service Principal ID.
- 4 Enter the API key.
- 5 Enter **SHA-256 fingerprint**. You can view and copy SHA-256 fingerprint from the NetBackup certificate authority details that is displayed on the NetBackup web UI.

See *View the NetBackup certificate authority details and fingerprint* topic in the *NetBackup Web UI Administrator's Guide*.

You can also view SHA-256 fingerprint by using command line. Run the following command on the NetBackup primary server:

```
/usr/opensv/netbackup/bin/nbcertcmd -listCACertDetails
```

See *NetBackup Commands Reference Guide*.

- 6 Click **Submit**.
- 7 In the **Ansible Output** tab, you can verify the details such as a new certificate on NetBackup OpenStack virtual machine that registers itself as a valid host on the NetBackup primary server.

Change NetBackup for OpenStack dashboard password

To change the NetBackup for OpenStack GUI password do:

- Log on to the NetBackup for OpenStack Dashboard.
- Click **Admin** in the upper right corner to open the submenu.

- Choose **Reset Password**.
- Set the new NetBackup for OpenStack password.

Reset NetBackup for OpenStack dashboard password

- Go to:
`/home/stack/myansible/lib/python3.6/site-packages/nbos_configurator/`
- Run: `/home/stack/myansible/bin/python recreate_conf.py`
- Restart **nbos-config** service: `systemctl restart nbos-config`

Downloading the NetBackup for OpenStack logs

You can download the NetBackup for OpenStack logs directly through the NetBackup for OpenStack configurator UI.

To download logs using the NetBackup for OpenStack configurator UI

- 1 Log on to the NetBackup for OpenStack configurator UI.
- 2 Go to **Logs**.
- 3 Select the logs to download.
 - Logs for every NetBackup for OpenStack appliance can be downloaded separately.
 - Zip of all log files can be created and downloaded.

This downloads the current log files. Already rotated logs need to be downloaded through SSH from the NetBackup for OpenStack appliance directly. All logs including rotated old logs can be found at the following location:

`/var/logs/nbosjm/`

Updating the API key

You can update the NetBackup service principle without having to perform an entire reconfiguration. This is especially useful if the service principle is expired, or the originally configured service principle is revoked.

To update the API key

- 1** Log on to the NetBackup for OpenStack configurator UI.
- 2** Click **API Key** at the left.
- 3** Under **API Key Validation**, update the service principal ID and the API key.
- 4** Click **Submit**.

Uploading the API certificate

You can upload the OpenStack cloud CA certificate where the OpenStack cloud is configured using HTTPS.

To update the API key

- 1** Log on to the NetBackup for OpenStack configurator UI.
- 2** Click **API Certificate** at the left.
- 3** Click **Certificate Authorities** to upload the certificate file.

Configuring NetBackup primary server

This chapter includes the following topics:

- [License for OpenStack plug-in for NetBackup](#)
- [About launching the OpenStack Horizon UI from the NetBackup web UI](#)
- [Configuring the NBOSVM service principal](#)
- [About NetBackup for OpenStack protection plan](#)
- [About auto image replication in NetBackup for OpenStack](#)

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

For more information on how to add licenses, see [NetBackup Administrator's Guide, Volume I](#)

About launching the OpenStack Horizon UI from the NetBackup web UI

You can access the Horizon UI by entering the horizon instance IP address or host name in the address bar.

You can also configure the Horizon instance details and launch the OpenStack Horizon UI from the NetBackup web UI.

Table 4-1 Launch OpenStack Horizon UI

Step	Task	Description
1	Add the OpenStack Horizon instance on the NetBackup web UI.	See “Adding the OpenStack Horizon instance on NetBackup web UI” on page 94.
2	Configure RBAC. <ul style="list-style-type: none"> Create a custom role for OpenStack administrator. Add users to a role. 	See “Creating the custom role for NetBackup for OpenStack administrator” on page 94.
3	Log on with the role, and launch the Horizon UI.	See “Launching the Horizon UI from the NetBackup web UI” on page 95.

Adding the OpenStack Horizon instance on NetBackup web UI

You can add the OpenStack Horizon instances on the NetBackup web UI and launch the Horizon UI from the web UI.

To add the OpenStack Horizon instances on the NetBackup web UI

- 1 On the web UI, click **OpenStack** under **Workload**.
- 2 Click **Add**.
- 3 In the Add Horizon instance link box, type the hostname/IP address and port number.
- 4 Click **Save**.

Creating the custom role for NetBackup for OpenStack administrator

The NetBackup web user interface provides the ability to apply role-based access control in your NetBackup environment. Use RBAC to provide access for the users that do not currently have access to NetBackup. Or, for current NetBackup users with administrator access you can provide limited access and permissions, based on their role in your organization.

For more information on configuring RBAC, see NetBackup™ web UI Administrator's Guide.

To add a custom role for NetBackup for OpenStack administrator

- 1 On the left, select **Security > RBAC**.
- 2 Select the **Roles** tab and click **Add**.
- 3 Select **Custom role** and click **Next**.

- 4 Provide a Role name and a description. For example, you may want to indicate that role is for any users that are backup administrators for a particular department or region.
- 5 For Role permissions, choose the permission or type of access that you want users with that role to have for each permission type.
- 6 Click **Add role**.

Launching the Horizon UI from the NetBackup web UI

After you create a custom role and add the users to the role, users with the custom role can log on to the Horizon UI.

To launch the Horizon UI from the NetBackup web UI

- 1 Sign in to the NetBackup web UI.
- 2 On the web UI, click **OpenStack** under **Workload**.
- 3 Click the URL.
- 4 Log on to the Horizon UI.

Configuring the NBOSVM service principal

You must configure service principal for a secure communication between NBOSVM and NetBackup.

Configuring the NBOSVM service principal

- 1 Create a non-root user in the NetBackup primary server.

```
adduser <username>
```
- 2 Log in to the NetBackup primary server web UI.
- 3 From the left side menu, go to **Security > RBAC > Default Security Administrator**.
- 4 On the **Users** tab, add the non-root user that you have created.
- 5 Go to **Security > Access keys**.

- 6 Click **Add** and enter the non-root user to create the access token.

7 Add the generated access token and `NetBackupHostName` in the cURL command and run it on the NetBackup primary server.

```
curl --insecure --location --request POST \

'https://<NetBackupHostName>:1556/netbackup/security/service-principal-configs' \
\
-H 'accept: application/vnd.netbackup+json;version=11.0' \
-H 'Content-Type: application/vnd.netbackup+json;version=11.0' \
\
-H 'Authorization: <Access Token>' \
-d '{
  "data": {
    "type": "servicePrincipalConfiguration",
    "attributes": {
      "servicePrincipalId": "Service_Principal_NBOSVM",
      "servicePrincipalType": "OPENSTACK",
      "servicePrincipalApiKeyExpireAfterDays": "P365D",
      "isSecurityAdmin": true,
      "accessDefinitions": [

        {
          "namespace": "|SECURITY|USERS|API-KEYS|",
          "operations": [

            "|OPERATIONS|VIEW|"
          ]
        },
        {
          "namespace": "|SECURITY|SERVICE-PRINCIPAL|",
          "operations": [

            "|OPERATIONS|VIEW|"
          ]
        },
        {
          "namespace": "|ASSETS|OPENSTACK|",
          "operations": [
            "|OPERATIONS|ADD|",
            "|OPERATIONS|VIEW|",
            "|OPERATIONS|UPDATE|",
            "|OPERATIONS|ASSETS|OPENSTACK|RESTORE_ORIGINAL|",
            "|OPERATIONS|ASSETS|OPENSTACK|RESTORE_ALTERNATE|",
```

```

        "|OPERATIONS|ASSETS|OPENSTACK|PROTECT|"
    ]
},
{
    "namespace": "|PROTECTION|PROTECTION_PLAN|",
    "operations": [
        "|OPERATIONS|VIEW|",
        "|OPERATIONS|PROTECTION|PROTECTION_PLAN|SUBSCRIBE|"
    ]
},
{
    "namespace": "|PROTECTION|POLICIES|",
    "operations": [
        "|OPERATIONS|PROTECTION|POLICIES|MANUAL-BACKUP|",
        "|OPERATIONS|VIEW|"
    ]
},
{
    "namespace": "|CREDENTIALS|",
    "operations": [

        "|OPERATIONS|ADD|",
        "|OPERATIONS|UPDATE|",
        "|OPERATIONS|DELETE|"
    ]
},
{
    "namespace": "|MANAGE|NBOSVM-SERVER|",
    "operations": [
        "|OPERATIONS|ADD|",
        "|OPERATIONS|UPDATE|",
        "|OPERATIONS|DELETE|"
    ]
},
{
    "namespace": "|MANAGE|JOBS|",
    "operations": [
        "|OPERATIONS|ADD|",
        "|OPERATIONS|VIEW|"
    ]
},
{
    "namespace": "|STORAGE|STORAGE-SERVERS|",

```

```

        "operations": [
            "|OPERATIONS|VIEW|"
        ]
    },
    {
        "namespace":
"|STORAGE|STORAGE-SERVERS|UNIVERSAL-SHARES|",
        "operations": [
            "|OPERATIONS|VIEW|"
        ]
    },
    {
        "namespace": "|MANAGE|IMAGES|",
        "operations": [
            "|OPERATIONS|VIEW|"
        ]
    }
]
}
}
} '

```

Note: Keep a note of `servicePrincipalId` and `apiKey` from the response of the cURL. They are required in the NetBackup for OpenStack configuration.

For information about `service-principal-configs` API, see the *NetBackup API Documentation*.

About NetBackup for OpenStack protection plan

You must create protection plan the NetBackup web UI.

For information on creating protection plans, see *NetBackup Web UI Administrator's Guide*.

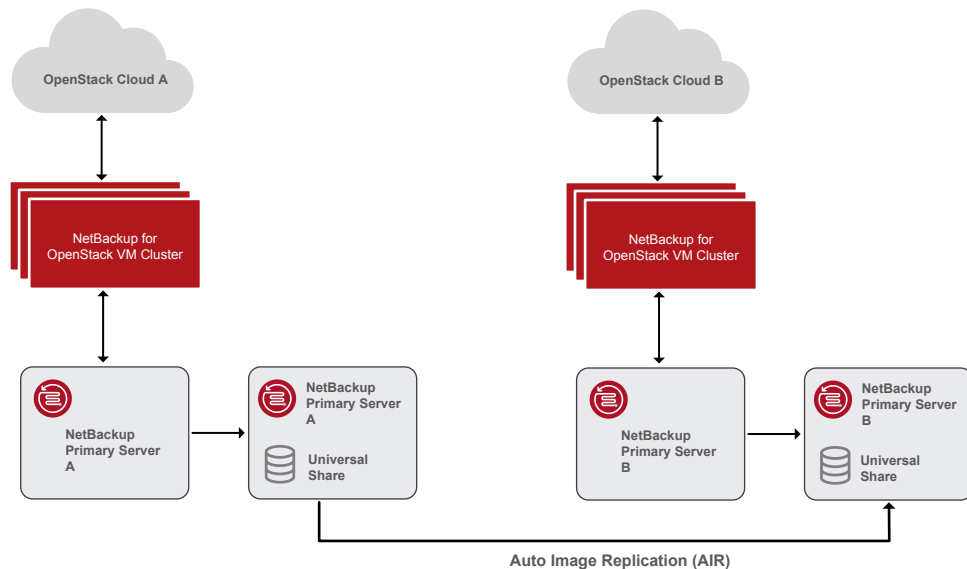
About auto image replication in NetBackup for OpenStack

The backups that are generated in one NetBackup domain can be replicated to storage in one or more target NetBackup domains. This process is referred to as auto image replication (AIR).

NetBackup for OpenStack 10.5 and later versions support AIR from a Media Server Deduplication Pool (MSDP) in one NetBackup domain to an MSDP in another domain. NetBackup uses storage lifecycle policies (SLP) in the source domain and the target domain to manage AIR operations.

The following diagram explains the architecture of using AIR between OpenStack Cloud A and OpenStack Cloud B. Backup images stored in the NetBackup media server A are replicated to the NetBackup media server B.

Figure 4-1 Auto image replication



See [“Configuring the AIR in NetBackup for OpenStack”](#) on page 100.

Configuring the AIR in NetBackup for OpenStack

After the AIR is configured between two NetBackup primary servers, the backup copy from the source NetBackup primary server is replicated to the target NetBackup primary server. You can import all the protections, recovery points, and the required metadata. After import of protections, you can recover the instance from the older backup copy.

To configure the AIR in NetBackup for OpenStack

- 1 Configure NBOSVM with source primary server using NBOSVM configurator UI.
- 2 Create a trusted relationship between source and target NetBackup primary servers.

*See [Configuring Auto Image Replication \(A.I.R.\)](#) topic of the *NetBackup Deduplication Guide*.*
- 3 On the source NetBackup primary server, add a replication target.

*See [Setting up the MSDP replication target](#) topic of the *NetBackup Deduplication Guide*.*
- 4 Create an Import SLP on the target NetBackup primary server.

*See [Configuring a storage lifecycle policy \(SLP\) for A.I.R.](#) topic of the *NetBackup Deduplication Guide*.*
- 5 Create a replication SLP on source NetBackup primary server.

*See [Configuring a storage lifecycle policy \(SLP\) for A.I.R.](#) topic of the *NetBackup Deduplication Guide*.*
- 6 Create a protection plan on source NetBackup primary server.
- 7 Create a protection on the Horizon UI with the protection plan.
- 8 Run a backup job.

To get the images on the target NBOSVM

- 1 On the target NBOSVM, import the protections from NetBackup by using the backup images stored on the universal share.

```
nbosjm protection-import-to-new-cloud
```

The protections are listed as orphaned protections. The project and user of the OpenStack cloud A do not exist on the OpenStack cloud B.

Note: After you import the protection in another cloud, when you enable the global job scheduler, all the protections that have scheduler trust enabled are displayed as broken because the protections do not have instances attached to them. Update the protection and assign an instance to it to enable the scheduler trust.

- 2 List the orphaned protections.

Orphaned protections are the protections that are no longer linked to an active tenant or user within the cloud. To identify and list all orphaned protections that

do not have `tenant_id` or `user_id` associated with the current cloud environment, run the following command:

```
nbosjm protection-get-orphaned-protections-list [--migrate_cloud
{True,False}]
```

- `--migrate_cloud` Set to `True` if you want to list policies from other clouds as well. The default value is `False`.

3 Assign protections to a new tenant or user.

```
nbosjm protection-reassign-protections [--old_tenant_ids
<old_tenant_id>]
                                     [--new_tenant_id
<new_tenant_id>]
                                     --protection_plan_id
<protection_plan_id>
                                     [--user_id <user_id>]
                                     [--migrate_cloud
{True,False}]
                                     [--map_file <map_file>]
```

- `--old_tenant_ids` The IDs of the old tenants from which you want to assign the protections.
- `--new_tenant_id` The ID of the new tenant to which you want to assign the protections.
- `--protection_plan_id` The ID of the protection plan to which you want to assign the protections.
- `--user_id` The user ID to which you want to assign the protections.
- `--migrate_cloud` Set to `True` if you want to assign the protections from other clouds also. The default value is `False`
- `--map_file` The file path with the file name of map file. The file format is `YAML`.

4 Run the following command to see the status of the import job.

```
nbosjm get-protection-import-status
```

NetBackup for OpenStack protections

This chapter includes the following topics:

- [About protections](#)
- [List of protections](#)
- [Create a protection](#)
- [Protection overview](#)
- [Edit a protection](#)
- [Delete a protection](#)
- [Unlock a protection](#)

About protections

A protection is a backup job that protects OpenStack virtual machines according to the configuration. You can create as many protections as needed, but each virtual machine can only be part of one protection.

List of protections

Using Horizon

To view all available protections of a project on Horizon

- ◆ On the Horizon console, navigate to **NBOS Backups > Protection**.

The overview in Horizon lists all the protections with the following additional information:

- Creation time
- Protection name
- Protection description
- Total recovery points inside this protection
 - Total number of succeeded recovery points
 - Total number of failed recovery points
- Protection description
- Protection type
- Protection status
- Scheduler Trust
 - Established denotes if the scheduler is enable for the protection.

Using CLI

```
nbosjm protection-list [--all {True,False}]
```

- `--all {True,False}` List all protections for all the projects. Valid for admin user only.

Create a protection

Using Horizon

To create a protection inside Horizon do the following steps:

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Click **Add protection**.
- 3 On the **Details** tab, provide the protection name, description, and the type as Serial or Parallel.
- 4 On the **Instances** tab, select the virtual machine to protect.
- 5 On the **Protection Plan** tab, select the protection plan from the drop-down list.

- 6 On the **Schedule** tab, Click **Enable Scheduler** to schedule the backups.
In the schedule, provide the start date, end date, start time, and the number of hours the snapshot/backup must repeat.
- 7 On the **Options** tab, you can pause the virtual machine during the snapshot creation. Select **Pause VM**.
- 8 Click **Create**.

The created protection will be available after a few seconds and starts to take backups according to the provided schedule and protection plan.

Using CLI

```
nbosjm protection-create [--protection-plan-id <protection plan_id>]
                        [--instance <instance-id=instance-uuid>]
                        [--display-name <display-name>]
                        [--display-description <display-description>]
                        [--protection-type-id <protection-type-id>]
                        [--source-platform <source-platform>]
                        [--jobschedule <key=key-name>]
                        [--metadata <key=key-name>]
```

- `--protection-plan-id` Protection plan ID to associate the protection with.
- `--display-name` The protection name.
- `--display-description` The protection description.
- `--protection-type-id` The protection type ID.
- `--source-platform` The protection source platform is required. openstack is the supported platform.
- `--instance` Specify an instance to include in the protection. Instance-id: Include the instance with this UUID.
- `--jobschedule` Specify the following key-value pairs for a job schedule. Specify the option multiple times to include multiple keys.

```
"start_date" : "06/05/2014"
"end_date"   : "07/15/2014"
"start_time" : "2:30 PM"
```

- `--metadata` Specify a key-value pair to include in the protection type metadata. Specify the option multiple times to include the multiple keys.

Protection overview

View the information about the protection in the protection overview.

Using Horizon

To enter the protection overview inside Horizon do the following steps:

- On the Horizon console, navigate to **NBOS Backups > Protection**.
- Identify the protection to view.
- Click the protection name to view the protection overview.

Details

The Protection Details tab provides the following information about the protection:

- Name
- Description
- Availability Zone
- Created at
- Last updated at
- Protection ID
- Protection type
- Protection plan name
- Protection plan ID
- Project ID
- User ID
- List of protected virtual machines including the information of qemu guest agent availability

The status of the QEMU guest agent shows whether the necessary OpenStack configuration is done for this virtual machine to provide QEMU guest agent integration. It does not check if the QEMU guest agent is installed and configured on the virtual machine.

You can navigate to the protected virtual machine directly from the list of protected virtual machines.

Recovery Point

The Recovery Point tab shows the list of all available recovery points in the chosen protection.

The copies are visible against the recovery points. These copies can be snapshot, backup, or duplicate copy.

See [“About recovery points”](#) on page 111.

Protection Plan

The Protection Plan tab gives an overview of the current configured scheduler and retention protection. The following elements are shown:

- Scheduler Enabled or Disabled
- Start Date and Time
- End Date and Time
- Repeat Every
- Time until the next snapshot runs
- Backup retention
- Backup retention
- Duplication

Using CLI

```
nbosjm protection-show <protection-id>
[--verbose <verbose>]
[--scheduler_trust <scheduler_trust {true}>]
```

- <protection-id> ID or name of the protection to show.
- --verbose Option to show additional information about the protection.
- --scheduler_trust Show protection for which schedule is enabled or not.

Edit a protection

A protection can be modified in all components to match changing needs.

Note: Editing a protection sets the user as the new owner.

Using Horizon

To edit a protection

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Identify the protection to be modified and select **Edit Protection** from the drop-down list.
- 3 Modify the protection as desired. All parameters except protection type can be changed.
- 4 Click **Update**.

Using CLI

```
nbosjm protection-modify [--display-name <display-name>]
                        [--display-description <display-description>]
                        [--instance <instance-id=instance-uuid>]
                        [--jobschedule <key=key-name>]
                        [--metadata <key=key-name>]
                        [--protection-plan-id <protection_plan_id>]
                        <protection-id>
```

- `--display-name` The protection name.
- `--display-description` The protection description.
- `--instance <instance-id=instance-uuid>` Specify an instance to include in the protection. Specify the option multiple times to include multiple instances. Instance-id: Include the instance with this UUID.

`"start_date" : "06/05/2014"`
`"end_date" : "07/15/2014"`
`"start_time" : "2:30 PM"`
- `--jobschedule <key=key-name>` Specify the following key-value pairs for a job schedule. Specify the option multiple times to include multiple keys. If you do not specify a time zone, it takes your local computer time zone by default.
- `--metadata` Specify a key-value pair to include in the protection type metadata. Specify the option multiple times to include multiple keys.
- `--protection-plan-id` ID of the protection.
- `<protection-id>` ID of the protection to edit.

Delete a protection

When a protection is no longer needed, you can delete it. You must expire all the recovery points before you delete the protection.

See [“About recovery points”](#) on page 111.

Using Horizon

To delete a protection

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Identify the protection to be modified and select **Delete Protection** from the drop-down list.
- 3 Click **Delete Protection** again to confirm.

Using CLI

```
nbosjm protection-delete [--database_only <True/False>] <protection-id>
```

- <protection-id> ID of the protection to delete.
- --database_only Keep True if you want to delete from the database only. The default value is False.

Unlock a protection

The protections that actively take backups or restores are locked for further tasks. You can unlock a protection by force if necessary.

Use this feature only as a last resort in case of backups or restores are stuck or a restore is required while a backup is running.

Using Horizon

To unlock a protection

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Identify the protection to be modified, and select **Unlock Protection** from the drop-down list.
- 3 Click **Unlock Protection** again to confirm.

Using CLI

```
nbosjm protection-unlock <protection-id>
```

- <protection-id> ID of the protection to unlock.

Performing snapshots, backups, and restores of OpenStack

This chapter includes the following topics:

- [About recovery points](#)
- [List of recovery points](#)
- [Creating a snapshot](#)
- [Snapshot and backup overview](#)
- [Expire recovery points](#)
- [Cleaning up the volume snapshots](#)
- [About restores](#)
- [List of Restores](#)
- [Restores overview](#)
- [Delete a restore](#)
- [Cancel a restore](#)
- [One-click restore](#)
- [Selective restore](#)
- [In-place restore](#)
- [Required restore.json file for CLI](#)

- [About backup mount](#)
- [Creating a file recovery manager instance](#)
- [Mounting a backup copy](#)
- [Accessing the file recovery manager](#)
- [Identifying mounted backups](#)
- [Unmounting a backup](#)
- [About schedules](#)
- [About activating the email notifications](#)

About recovery points

A recovery point is a single NetBackup for OpenStack backup of a protection including all data and metadata. It contains the information of all virtual machines that the protection protects.

List of recovery points

Using Horizon

To view the list of the recovery points

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Identify the protection to show the details on.
- 3 Click the protection name to enter the protection overview.
- 4 Navigate to the **Recovery Points** tab.

The list of recovery points for the chosen protection contains the following additional information:

- Creation Time
- Name of the recovery point
- Description of the recovery point
- The number of restores from this recovery point
 - The number of succeeded Restores
 - The number of failed Restores

- Status
- Copies
- Action

Using CLI

```
nbosjm recovery-point-list [--protection-id <protection-id>]
                           [--nbos_node <host>]
                           [--date_from <date_from>]
                           [--date_to <date_to>]
                           [--all {True,False}]
```

- `--protection-id` Filter results by protection-id (protection ID).
- `--nbos_node` List all the recovery points operations that are scheduled on a NetBackup for OpenStack node. The default value is None.
- `--date_from` From date in the format YYYY-MM-DDTHH:MM:SS. For example, 2016-10-10T00:00:00. If you do not specify the time, it takes 00:00 by default.
- `--date_to` To date in the format YYYY-MM-DDTHH:MM:SS The default is current day. Specify the time in HH:MM:SS format to get the recovery points in the same day.
- `--all` List all recovery points of all the projects. Valid for admin user only.

Creating a snapshot

Snapshots are automatically created by the NetBackup for OpenStack scheduler. If necessary or in the case of a deactivated scheduler, you can create a snapshot on demand.

Note: NetBackup for OpenStack does not support backup of swap disks and ephemeral disks.

Using Horizon

You can create a snapshot from the protection overview and the protection snapshot list page.

To create a snapshot from the protection overview

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Identify the protection to create a snapshot.

- 3 Click **Backup Now** to create a snapshot.
- 4 Provide a name and description for the snapshot.
- 5 Click **Create**.

To create a snapshot from the protection snapshot

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Identify the protection to create a snapshot.
- 3 Click the protection name to enter the protection overview.
- 4 On the **Recovery Points** tab, click **Backup Now**.
- 5 Provide a name and description for the snapshot.
- 6 Click **Create**.

Using CLI

```
nbosjm protection-snapshot [--display-name <display-name>]
                             [--display-description <display-description>]
                             <protection-id>
```

- <protection-id> ID of the protection to create a snapshot.
- --display-name The name of the snapshot.
- --display-description The snapshot description.

Snapshot and backup overview

Each recovery point contains information about the snapshot and the backup copies. This information can be seen in the recovery point overview.

If a nova-booted instance is migrated from one compute node to another after the snapshot operation is completed, backup or restore from the snapshot copy is not supported.

If a backup copy is mounted on a file recovery manager instance, you must unmount the backup copy to take the backup of the file recovery manager instance.

Using Horizon

To reach the recovery point overview follow these steps:

1. On the Horizon console, navigate to **NBOS Backups > Protection**.
2. Identify the protection that contains the recovery point to show.

3. Click the protection name to enter the protection overview.
4. Navigate to the **Recovery Points** tab.
5. Identify the searched recovery point in the recovery point list.
6. Click the recovery point name.

Details	<p>The Recovery Points Details tab shows the following information about the recovery point.</p> <ul style="list-style-type: none"> ■ ID/Name/Description ■ Scheduled on ■ Total volume size ■ Snapshot Type ■ Snapshot Size ■ Snapshot Time Taken ■ Snapshot Status ■ Backup Size ■ Backup Time Taken ■ Backup Status ■ Backup Type ■ Virtual machines that are part of the recovery point ■ The following information is displayed for each virtual machine in the recovery point: <ul style="list-style-type: none"> ■ Instance Info - Name and Status ■ Security Group(s) - Name, Type ■ Flavor - vCPUs, Disk, and RAM ■ Networks - IP, Network name, and Mac Address ■ Attached Volumes - Name, Type, Size (GB), and Device Path ■ Misc - Original ID of the virtual machine
Restores	<p>The Restores tab shows the list of restores that have been started from the chosen recovery point. You can start the restores from here.</p> <p>See “About restores” on page 116.</p>
Misc.	<p>The Miscellaneous tab provides the remaining metadata information about the snapshot.</p> <ul style="list-style-type: none"> ■ Creation Time ■ Last Update time ■ ID ■ Protection ID of the protection containing the snapshot

Using CLI

```
nbosjm recovery-point-show [--output <output>] <recovery_point_id>
```

- `<recovery_point_id>` ID of the recovery point to be shown.
- `--output <output>` Option to get additional snapshot details.
 - Specify `--output metadata` for recovery point metadata.
 - Specify `--output networks` for snapshot virtual machines networks.
 - Specify `--output disks` for snapshot virtual machines disks.
 - Specify `--output copies` to list all the copies (snapshot, backup, and duplicate copies).

Note: OpenStack does not let you launch an instance without a network interface. The snapshot of the instance that does not have any network interface attached to it cannot be restored using the selective restore or one-click restore options. However, you can use in-place restore, which does not launch an instance.

Expire recovery points

When a recovery point is expired, an image cleanup operation is triggered from the primary server. This operation also cleans up the volume snapshots, which are part of the recovery point.

Cleaning up the volume snapshots

You can clean up the volume snapshots that are in the failed state or error state using the snapshot ID or the protection ID.

Using CLI

```
nbosjm volume-snapshot-cleanup --recovery_point_id <recovery_point_id>  
nbosjm volume-snapshot-cleanup --protection_id <protection_id>
```

- `<recovery_point_id>` ID of the recovery point for which volume snapshot cleanup is performed.
- `<protection_id>` ID of the protection for which volume snapshot cleanup is performed.

Note: If you use both `recovery_point_id` and `protection_id` options, snapshot ID must be associated with the protection.

About restores

A Restore is the workflow to bring back the backed-up virtual machines from the NetBackup for OpenStack snapshot, backup, or duplicate copies.

Note: If a nova-booted instance is migrated from one compute node to another after the snapshot operation is completed, backup or restore from the snapshot is not supported.

About restoring the multi-attach volumes

NetBackup for OpenStack supports multi-attach volumes for backup and restore. With this feature, you can share one volume with the multiple virtual machines. For more information about multi-attach volumes, see the *OpenStack documentation*.

During the backup of virtual machines with multi-attach volumes, each virtual machine is backed up separately. Hence, when restore operations on virtual machines with multi-attach volumes is performed, the restored volume has a multi-attach property set but is not shared by default.

For example, you have a multi-attach volume that is attached to four virtual machines protected by four different protections. You take the backup of four virtual machines with these protections. When you restore the instance from the snapshot or backup copy, it restores four virtual machines with four separate volumes with the multi-attach property set.

List of Restores

Using Horizon

To reach the list of Restores for a recovery point follow these steps:

1. On the Horizon console, navigate to **NBOS Backups > Protection**.
2. Identify the protection that contains the recovery point to show.
3. Click the protection name to enter the protection overview.
4. Navigate to the **Recovery Points** tab.
5. Identify the recovery point in the recovery points list.

6. Click the recovery point name.
7. Navigate to the **Restores** tab.

Using CLI

```
nbosjm restore-list [--recovery_point_id <recovery_point_id>]
```

- `--recovery_point_id` Filter the results by an ID of the recovery point.

Restores overview

Using Horizon

To reach the detailed Restore overview follow these steps:

1. On the Horizon console, navigate to **NBOS Backups > Protection**.
2. Identify the protection that contains the snapshot to show.
3. Click the protection name to enter the protection overview.
4. Navigate to the **Recovery Points** tab.
5. Identify the recovery point in the recovery point list.
6. Click the recovery point name.
7. Navigate to the **Restores** tab.
8. Identify the restore to show.
9. Click the restore name.

Details

The Restore Details Tab shows the following information about the restore:

- Name
- Description
- Restore Type
- Status
- Time taken
- Size
- Progress Message
- Progress
- Host
- Restore Options

The Restore Options are the `restore.json` provided to NetBackup for OpenStack.

- List of virtual machines restored
 - Restored virtual machine name
 - Restored virtual machine status
 - Restored virtual machine ID
- NetBackup Copy Number
- NetBackup Copy Type

Misc

The Misc tab provides additional metadata information.

- Creation Time
- Restore ID
- Recovery point ID containing the restore
- Protection

Using CLI

```
nbosjm restore-show [--output <output>] <restore_id>
```

- `<restore_id>` ID of the restore to be shown.
- `--output <output>` Option to get additional restore details.

Specify `-output metadata` for restore metadata

```
-output networks
```

```
-output subnets
```

```
-output routers
```

```
-output flavors
```

Delete a restore

Once a restore is no longer needed, it can be safely deleted from a protection.

Deleting a restore only deletes the NetBackup for OpenStack information about this restore. No OpenStack resources are deleted.

Using Horizon

Deleting a single restore through the submenu

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Identify the protection that contains the recovery point to delete.
- 3 Click the protection name to enter the protection overview.
- 4 Navigate to the **Recovery Points** tab.
- 5 Identify the searched recovery point in the recovery point list.
- 6 Click the recovery point name.
- 7 Navigate to the **Restore** tab.
- 8 Click **Delete Restore** in the line of the restore in question.
- 9 Click **Delete Restore** again to confirm.

Deleting the multiple restores through a checkbox in recovery point overview

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Identify the protection that contains the recovery point to show.
- 3 Click the protection name to enter the protection overview.
- 4 Navigate to the **Recovery Points** tab.
- 5 Identify the searched recovery point in the recovery points list.
- 6 Enter the recovery point by clicking the recovery point name.
- 7 Navigate to the **Restore** tab.
- 8 Select the check box for each restore that shall be deleted.
- 9 Click **Delete Restore**.
- 10 Click **Delete Restore** again to confirm.

Using CLI

```
nbosjm restore-delete <restores_id>
```

- <restore_id> ID of the restore to be deleted.

Cancel a restore

Ongoing restores can be canceled using the command line.

Using CLI

```
nbosjm restore-cancel <restore_id>
```

- <restore_id> ID of the restore to be deleted.

One-click restore

The one-click restore brings back all virtual machines from the snapshot or backup in the same state as they were backed up. They are located in the same cluster in the same datacenter, use the same storage domain, connect to the same network, and have the same flavor.

The user cannot change any metadata.

The one-click restore requires that the original virtual machines that have been backed up are deleted or otherwise lost. Even if one virtual machine is still running, the one-click restore fails.

The one-click restore automatically updates the protection to protect the restored virtual machines.

Note: One-click restore fails if the properties of the instances that existed at the time of snapshot or backup do not exist at the time of restore.

Using Horizon

To perform the one-click restore

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Identify the protection that contains the recovery point to be restored.
- 3 Click the protection name to enter the protection overview.
- 4 Navigate to the **Recovery points** tab.
- 5 Identify the recovery point to be restored.
- 6 Click **One-Click Restore** in the same line as the identified recovery point.
- 7 (Optional) Provide the name and description.
- 8 Click **Create**.

Using CLI

```
nbosjm oneclick-restore [--display-name <display-name>]
                        [--display-description <display-description>]
                        <recovery_point_id>
                        <copy_number>
                        <copy_type>
```

- `<recovery_point_id>` ID of the recovery point to restore.
- `<copy_number>` Copy number of snapshot or backup for restore.
- `<copy_type>` Specify copy type of snapshot or backup for the restore.
- `--display-name` Optional name for the restore.
- `--display-description` Optional description for restore.

Selective restore

The selective restore is the most complex restore NetBackup for OpenStack has to offer. It allows to adapt the restored virtual machines to the exact needs of the user.

With the selective restore the following things can be changed:

- Which virtual machines are getting restored.
- Name of the restored virtual machines
- Which networks to connect with.
- Which Storage domain to use
- Which datacenter or cluster to restore into.
- Which flavor the restored virtual machines will use.

The selective restore is always available and does not have any prerequisites.

Using Horizon

To perform a selective restore

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Identify the protection that contains the recovery point to be restored.
- 3 Click the protection name to enter the protection overview.
- 4 Navigate to the **Recovery points** tab.

- 5 Identify the recovery point to be restored.
- 6 From the dropdown menu under the **Actions** column, select **Selective Restore**.
- 7 Configure the selective restore as desired.
- 8 Click **Restore**.

Using CLI

```
nbosjm selective-restore [--display-name <display-name>]
                        [--display-description <display-description>]
                        [--filename <filename>]
                        <recovery_point_id>
```

- <recovery_point_id> ID of the recovery point to restore.
- --display-name Optional name for the restore.
- --display-description Optional description for restore.
- --filename Provide the file path (relative or absolute) including the file name.
By default it reads the file `/home/stack/myansible/lib/python3.8/site-packages/nbosjmclient/input-files/restore_from_backup_copy.json`.
You can use this file for a reference or replace values into this file.

In-place restore

The in-place restore covers those use cases, where the virtual machine and its volumes are still available, but the data got corrupted or needs to rollback for other reasons.

It allows the user to restore only the data of a selected volume, which is part of a backup.

The in-place restore only works when the original virtual machine and the original volume are still available and connected. NetBackup for OpenStack checks the status with the saved Object-ID.

The in-place restore will not create any new RHV resources. Use one of the other restore options if new volumes or virtual machines are required.

The in-place restore restarts the instance.

Note: The in-place restore does not support a restore from the snapshot.

Note: In-place restore fails if the properties of the instances that existed at the time of backup do not exist at the time of restore.

Using Horizon

To perform the in-place restore

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Identify the protection that contains the recovery point to be restored.
- 3 Click the protection name to enter the protection overview.
- 4 Navigate to the **Recovery Points** tab.
- 5 Identify the recovery point to be restored.
- 6 From the drop-down under **Actions** column, select **Inplace Restore**.
- 7 Configure the In-place restore as desired.
- 8 Click **Restore**.

Using CLI

```
nbosjm inplace-restore [--display-name <display-name>]
                        [--display-description <display-description>]
                        [--filename <filename>]
                        <recovery_point_id>
```

- <recovery_point_id> ID of the recovery point to restore from the backup copy.
- --display-name Optional name for the restore.
- --display-description Optional description for restore.
- --filename Provide file path (relative or absolute) including file name. By default it reads the file `/home/stack/myansible/lib/python3.8/site-packages/nbosjmcclient/input-files/restore_from_backup_copy.json`. You can use this file for reference or replace values into this file.

Required restore.json file for CLI

The nbosjm client CLI uses a `restore.json` file to define the restore parameters for the selective and the in-place restore.

An example for a selective restore of this `restore.json` is shown below. A detailed analysis and explanation is given afterwards.

The restore.json requires information about the backed-up resources. All required information can be gathered in the recovery point overview.

```
{
  'name': 'sel-rest-5',
  'description': 'sel-rest-desc-5',
  'oneclickrestore': False,
  'restore_type': 'selective',
  'copy_number': '2',
  'copy_type': 'BACKUP',
  'type': 'openstack',
  'openstack':
    {
      'restore_topology': False,
      'instances':
        [
          {
            'id': '91a26084-7134-4ae4-970c-8203fb18669f',
            'name': 'sample-instance-restore',
            'restore_boot_disk': True,
            'availability_zone': 'nova',
            'include': True,
            'vdisks':
              [
                {
                  'id': 'c6fe8309-a95b-4bbb-9d72-57beafe4a3ae',
                  'new_volume_type': '__DEFAULT__',
                  'availability_zone': 'nova'
                }
              ]
          },
          {
            'id': 'd8680981-2113-45a8-aa7c-6edd68c97819',
            'mac_address': 'fa:16:3e:d1:ce:ae',
            'network': {
              'id': 'd8680981-2113-45a8-aa7c-6edd68c97819',
              'subnet': {
                'id': '28206b2e-0a0e-46a3-9034-9d621b4bfb4f'
              }
            },
            'ip_address': '172.20.2.230'
          }
        ]
    }
}
```

```

    }
  ],
  'networks_mapping': {
    'networks': [
      { 'snapshot_network': {
        'name': 'private',
        'id': 'd8680981-2113-45a8-aa7c-6edd68c97819',
        'subnet': {
          'id': '28206b2e-0a0e-46a3-9034-9d621b4bfb4f'
        }
      } },
      { 'target_network': {
        'id': 'd8680981-2113-45a8-aa7c-6edd68c97819',
        'name': 'private',
        'subnet': {
          'id': '28206b2e-0a0e-46a3-9034-9d621b4bfb4f'
        }
      } }
    ]
  }
}

```

General required information

Before the exact details of the restore are to be provided it is necessary to provide the general metadata for the restore.

- `name` The name of the restore.
- `description` The description of the restore.
- `oneclickrestore <True/False>` If the restore is a one-click restore. Setting this option to True will override all other settings and a one-click restore is started.
- `restore_type <oneclick/selective/inplace>` Defines the restore that is intended .
- `type openstack` Defines that the restore is into an OpenStack cloud.
- `openstack` Starts the exact definition of the restore.
- `copy number` The number of the backup copy.
- `copy_type` The format of the data.

Selective restore required information

The selective restore requires a lot of information to be able to execute the restore as desired.

The information is divided into three components:

- `instances`
- `restore_topology`
- `networks_mapping`

Information required in instances

This part contains all information about all instances that are part of the recovery point to restore and how they are to be restored.

Even when virtual machines are not to be restored, they are required to be inside the `restore.json` to allow a clean execution of the restore.

Each instance requires the following information.

- `id` Original ID of the instance.
- `include` <True/False> Set True when the instance shall be restored.

All further information is only required, when the instance is part of the restore.

- `name` New name of the instance.
- `availability_zone` Nova Availability Zone the instance shall be restored into. Leave empty for "Any Availability Zone".
- `Nics` List of the OpenStack Neutron ports that shall be attached to the instance. Each Neutron Port consists of:
 - `id` ID of the Neutron port to use
 - `mac_address` Mac Address of the Neutron port
 - `ip_address` IP address of the Neutron port
 - `network` Network the port is assigned to. Contains the following information:
 - `id` ID of the network the Neutron port is part of.
 - `subnet` Subnet the port is assigned to. Contains the following information:
 - `id` ID of the network the Neutron port is part of.

To use the next free IP available, set `Nics` to an empty list []

Using an empty list for `Nics` combined with the network topology restore, the restore job sets the original IP address of the instance.

- **vdisks** List of all volumes that are part of the instance. Each volume requires the following information:
 - **id** Original ID of the volume.
 - **new_volume_type** The volume type to use for the restored volume. Leave empty for Volume Type None.
 - **availability_zone** The Cinder Availability Zone to use for the volume. The default Availability Zone of Cinder is Nova.
- **flavor** Defines the Flavor to use for the restored instance. Contains the following information:
 - **ram** How much RAM the restored instance will have (in MB).
 - **ephemeral** How big the ephemeral disk of the instance will be (in GB).
 - **vcpus** How many vcpus the restored instance will have available.
 - **swap** How big the Swap of the restored instance will be (in MB). Leave empty for none.
 - **disk** Size of the root disk the instance will start with.
 - **id** ID of the flavor that matches the provided information.

Warning: The root disk needs to be at least as big as the root disk of the backed-up instance.

The following example describes a single instance with all values.

```
'instances':[
  {
    'name':'cdcentOS-1-selective',
    'availability_zone':'US-East',
    'nics':[
      {
        'mac_address':'fa:16:3e:00:bd:60',
        'ip_address':'192.168.0.100',
        'id':'8b871820-f92e-41f6-80b4-00555a649b4c',
        'network':{
          'subnet':{
            'id':'2b1506f4-2a7a-4602-a8b9-b7e8a49f95b8'
          },
          'id':'d5047e84-077e-4b38-bc43-e3360b0ad174'
        }
      }
    ]
  }
]
```

```

    }
  ],
  'vdisks':[
    {
      'id':'4cc2b474-1f1b-4054-a922-497ef5564624',
      'new_volume_type':'ceph',
      'availability_zone':'nova'
    }
  ],
  'flavor':{
    'ram':2048,
    'ephemeral':0,
    'vcpus':1,
    'swap':'',
    'disk':20,
    'id':'2'
  },
  'include':True,
  'id':'890888bc-a001-4b62-a25b-484b34ac6e7e'
}
]

```

Information required in network topology restore or network mapping

Warning: Do not mix network topology restore together with network mapping.

To activate a network topology restore set:

```
restore_topology:True
```

To activate network-mapping set:

```
restore_topology:False
```

When the network mapping is activated it is used, it is necessary to provide the mapping details, which are part of the networks_mapping block:

- `networks` List of snapshot_network and target_network pairs.
 - `snapshot_network` The network backed up in the snapshot, contains the following:

- `id` Original ID of the network backed up.
- `subnet` The subnet of the network that is backed up in the snapshot, contains the following:
 - `id` Original ID of the subnet backed up.
- `target_network` The existing network to map to, contains the following:
 - `id` ID of the network to map to.
 - `subnet` The subnet of the network backed up in the snapshot, contains the following:
 - `id` ID of the subnet to map to.

Full selective restore example

```
{
  'description': 'u  -',
  'oneclickrestore': False,
  'openstack': {
    'instances': [
      {
        'name': 'cdcentOS-1-selective',
        'availability_zone': 'US-East',
        'nics': [
          {
            'mac_address': 'fa:16:3e:00:bd:60',
            'ip_address': '192.168.0.100',
            'id': '8b871820-f92e-41f6-80b4-00555a649b4c',
            'network': {
              'subnet': {
                'id': '2b1506f4-2a7a-4602-a8b9-b7e8a49f95b8'
              },
              'id': 'd5047e84-077e-4b38-bc43-e3360b0ad174'
            }
          }
        ],
        'vdisks': [
          {
            'id': '4cc2b474-1f1b-4054-a922-497ef5564624',
            'new_volume_type': 'ceph',
            'availability_zone': 'nova'
          }
        ]
      }
    ]
  }
}
```

```

],
  'flavor':{
    'ram':2048,
    'ephemeral':0,
    'vcpus':1,
    'swap':'',
    'disk':20,
    'id':'2'
  },
  'include':True,
  'id':'890888bc-a001-4b62-a25b-484b34ac6e7e'
}
],
'restore_topology':False,
'networks_mapping':{
  'networks':[
    {
      'snapshot_network':{
        'subnet':{
          'id':'8b609440-4abf-4acf-a36b-9a0fa70c383c'
        },
        'id':'8b871820-f92e-41f6-80b4-00555a649b4c'
      },
      'target_network':{
        'subnet':{
          'id':'2b1506f4-2a7a-4602-a8b9-b7e8a49f95b8'
        },
        'id':'d5047e84-077e-4b38-bc43-e3360b0ad174',
        'name':'internal'
      }
    }
  ]
}
},
'restore_type':'selective',
'type':'openstack',
'name':'getjson2'
}

```

In-place restore required information

The in-place restore requires less information than a selective restore. It only requires the base file with some information about the instances and volumes to be restored.

Information required in instances

- `id` ID of the instance inside the Snapshot.
- `restore_boot_disk` Set to True if the boot disk of that virtual machine shall be restored.

When the boot disk is at the same time a Cinder Disk, both values need to be set true.

- `include` Set to True if at least one volume from this instance shall be restored.
- `vdisk` List of the disks that are connected to the instance. Each disk contains:
 - `id` Original ID of the volume.
 - `restore_cinder_volume` Set to true if the volume shall be restored.
 - `new_volume_type` Volume type of the restored volume. Set to the same value as the original volume.

Network mapping information required

There is no network information required, but the field has to exist as an empty value for the restore to work.

Full in-place restore example

```
{
  'description':u    '-',
  'name':'Inplace Restore',
  'zone':'',
  'oneclickrestore':False,
  'restore_type':u    'inplace',
  'type':u    'openstack',
  'openstack':{
    'instances':[
      {
        'restore_boot_disk':True,
        'include':True,
        'id':'ba8c27ab-06ed-4451-9922-d919171078de',
        'vdisk':[
          {
            'restore_cinder_volume':True,
            'id':'04d66b70-6d7c-4d1b-98e0-11059b89cba6',
            'new_volume_type':'ceph'
          }
        ]
      }
    ]
  }
}
```

```
        }  
      ]  
    }  
  ],  
  'restore_topology':False,  
  'networks_mapping':{  
    'networks':[  
    ]  
  }  
}  
}
```

About backup mount

NetBackup for OpenStack lets you view or download a file from the backup copy. Any changes to the files or directories when backup copy is mounted are temporary and are discarded when the backup copy is unmounted. Mounting is a faster way to restore a single or multiple files. To mount a backup copy follow these steps.

Creating a file recovery manager instance

Create an OpenStack image using a Linux-based cloud image like Ubuntu or RHEL 8.2 or later. Add the following metadata parameters and upload the cloud image to the Glance.

```
--file <File Manager Image Path> \  
--container-format bare \  
--disk-format qcow2 \  
--public \  
--property hw_gemu_guest_agent=yes \  
--property nbos_recovery_manager=yes \  
--property hw_disk_bus=virtio \  
nbos-file-manager
```

Spin up an instance from that image. It is recommended to have at least 8GB RAM for the mount operation. A large backup copy may require more RAM.

Steps to apply on RHEL 8.2 or later cloud images

- 1 Install and activate **qemu-guest-agent**.
- 2 Edit `/etc/sysconfig/qemu-ga` and remove the following from the `BLACKLIST_RPC` section.

```
guest-file-read
guest-file-write
guest-file-open
guest-file-close
```

- 3 Disable SELINUX in `/etc/sysconfig/selinux`.

```
SELINUX=disabled
```

- 4 Install python3.

```
yum install python3
```

- 5 Install lvm2.

```
yum install lvm2
```

- 6 Restart the instance.

Steps to apply on Ubuntu cloud images

- 1 Install and activate **qemu-guest-agent**.
- 2 Edit `/etc/init.d/qemu-guest-agent` and add Freeze-Hook file path in daemon args.

```
DAEMON_ARGS="-F /etc/qemu/fsfreeze-hook"
```

- 3 Generate `qemu-ga.conf` file.

```
qemu-ga -D > /etc/qemu/qemu-ga.conf
```

- 4 Append following line to the file.

```
fsfreeze-hook=/etc/qemu/fsfreeze-hook
```

- 5 Restart the `qemu-guest-agent` service.

- 6 Install Python 3.

```
apt-get install python3
```

- 7 Restart the instance.

Mounting a backup copy

Mounting a backup copy to a file recovery manager instance provides read access to all the data that is located in the mounted backup copy.

Unmount any mounted backup when there is no further need to keep it mounted. The retention policy does not purge the mounted backups.

You can run the mounting process against any OpenStack instance. During this process the instance is restarted.

During the mounting process, the OpenStack instance is restarted.

Always mount backups to file recovery manager instances only.

Note: Mirrored volumes are not mounted automatically on the file recovery manager instance. You must mount the mirrored volumes manually.

To mount a backup copy

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Identify the protection that contains the backup to mount.
- 3 Click the protection name to enter the protection overview.
- 4 Navigate to the **Recovery Points** tab.
- 5 Click **Copies** at the right of the recovery point row.
- 6 Identify the backup copy and select **OneClick Restore** from the drop-down list.
- 7 Click **Mount for file restore**.
- 8 Choose the file recovery manager instance to mount to.
- 9 Click **Mount** to confirm.

If all instances of the project are listed and there is a file recovery manager instance, verify that the file recovery manager image has the following property set:

```
nbos_recovery_manager=yes
```

Using CLI

```
nbosjm backup-mount <mount_vm_id> <copy_number> <recovery_point_id>
```

- <mount_vm_id> VM ID that backup volumes mount to.
- <copy_number> Specify copy number for backup mount.

- `<recovery_point_id>` ID of the recovery point to mount.

Accessing the file recovery manager

The file recovery manager is a normal Linux based OpenStack instance.

It can be accessed by SSH or SSH-based tools like FileZilla or WinSCP.

SSH login is often disabled by default in cloud-images. Enable SSH login if necessary.

The mounted backup copy can be found at the following path:

```
/home/ubuntu/nbos-mounts/mounts/
```

Each virtual machine has its own directory using the VM_ID as the identifier.

Identifying mounted backups

Sometimes backup copy is mounted for a longer duration and hence it is important to be identified.

Using Horizon

There are 2 possibilities to identify mounted backups inside Horizon.

From the file recovery manager instance metadata

1. On the Horizon console, navigate to **Compute > Instances**.
2. Identify the file recovery manager instance.
3. Click the name of the file recovery manager instance to bring up its details.
4. On the **Overview** tab look for Metadata.
5. Identify the value for `mounted_snapshot_url`

The `mounted_snapshot_url` contains the ID of the backup that has been mounted last.

Note: This value only gets updated, when a new backup is mounted.

From the recovery points list

1. On the Horizon console, navigate to **NBOS Backups > Protection**.
2. Identify the protection that contains the backup to mount.
3. Click the protection name to enter the protection overview.

4. Navigate to the **Recovery Points** tab.
5. Click **Copies** at the right of the recovery point row.
6. Search for the backup copy that has the option **Unmount Backup**.

Using CLI

```
nbosjm backup-mounted-list
```

- List of all mounted backups.

Unmounting a backup

Once a mounted backup is no longer needed it is possible and recommended to unmount the backup.

Unmounting a backup frees the file recovery manager instance to mount the next backup and allows NetBackup for OpenStack retention policy to purge the former mounted backup.

Warning: Deleting the file recovery manager instance does not update the NetBackup for OpenStack appliance. The backup will be considered mounted until an unmount command has been received.

Using Horizon

1. On the Horizon console, navigate to **NBOS Backups > Protection**.
2. Identify the protection that contains the backup to unmount.
3. Click the protection name to enter the protection overview.
4. Navigate to the **Recovery Points** tab.
5. Click **Copies** at the right of the recovery point row.
6. Identify the backup copy and click **Unmount Backup**.

Using CLI

```
nbosjm backup-dismount <recovery_point_id>
```

- <recovery_point_id> ID of the recovery point to dismount.

About schedules

Every protection has its own schedule. Those schedules can be activated, deactivated, and modified.

A schedule is defined by:

- Status (Enabled/Disabled)
- Start Day/Time
- End Day
- Hrs between two snapshots

Enabling or disabling a schedule

You can enable or disable the scheduler of a single protection using Horizon and command-line interfaces.

To enable or disable the scheduler of a single protection using Horizon

- 1 On the Horizon console, navigate to **NBOS Backups > Protection**.
- 2 Identify the protection to be modified.
- 3 From the drop-down under **Actions** column, select **Edit Protection**.
- 4 Navigate to the tab **Schedule**.
- 5 Select **Enabled** or **Disabled**.
- 6 Click **Update**.

To enable or disable the scheduler of a single protection using the command-line

- 1 Run the following command to enable the scheduler.

```
nbosjm enable-scheduler --protectionids <protectionid>
```
- 2 Run the following command to disable the scheduler.

```
nbosjm disable-scheduler --protectionids <protectionid>
```

 - `--protectionids` Requires at least one protection ID. Specify an ID of the protection to enable or disable the schedule for. Specify an option multiple times to include multiple policies.

Modifying a schedule

To modify a schedule, you must modify the protection.

See [“Edit a protection”](#) on page 107.

About activating the email notifications

NetBackup for OpenStack sends email notifications after every backup and restore. The email is sent to the owner of the protection.

The OpenStack administrator must ensure that the following requirements are met to activate the email notifications:

- User email is assigned
As the email is sent to the owner of the protection, the OpenStack user who created the protection is required to have an email address associated.
- NetBackup for OpenStack mail server is configured
NetBackup for OpenStack needs to know which mail server to use to send the email notifications. The backup administrators can configure mail server on Horizon.

To activate the email notifications

- 1 On the Horizon console, navigate to **NBOS Backups > Settings**.
- 2 Select or clear the box for **Enable Email Alerts**.

Performing Backup Administration tasks

This chapter includes the following topics:

- [NBOS Backup Admin Area](#)
- [Protection plan](#)
- [Managing the trusts](#)
- [Policy import and migration](#)

NBOS Backup Admin Area

NetBackup for OpenStack provides Backup as a Service, which allows OpenStack users to manage and control their backups themselves.

To provide backup administrators with the tools they need, NetBackup for OpenStack provides NBOS Backup Admin area in Horizon in addition to the API and the command-line interface.

Access the NBOS Backup Admin area

Access the NBOS Backup Admin area

- 1 Log in to the Horizon console with the administrator user.
- 2 Navigate to **Admin > NBOS Backup Admin > NetBackupOpenStack**.

You can filter and view the information for a specific tenant also.

Status overview

The status overview is always visible in the NBOS Backup Admin area. It provides the following information:

- Number of protected virtual machines in comparison with the number of existing virtual machines
- Number of currently running snapshots
- Status of NBOS nodes
- Status of NBOSDM services

The status of nodes is filled when the services are running and in good status.

Subscriptions tab

This tab provides information about all currently existing protections. It is the most important overview tab for every backup administrator and therefore the default tab is shown when the NBOS Backup Admin area is opened.

The following information is shown:

- User-ID that owns the protection.
- Project that contains the protection.
- Subscription
- Protection type
- Availability zone
- Number of protected virtual machines
- Performance information about the last 30 backups
 - How much data was backed up (green bars)
 - How long did the backup take (red line)
- Pie chart that shows the number of full backups and incremental backups.
- Number of successful backups
- Number of failed backups
- Storage that is used by that protection.
- Which backup target is used.
- When is the next snapshot run.
- What is the general interval of the protection
- Scheduler status including a switch to deactivate/activate the protection

Usage tab

Administrators often need to figure out where a lot of resources are used or they need to quickly provide usage information to a billing system. This tab helps in these tasks by providing the following information:

- Storage used by a tenant
- Virtual machines protected by a tenant

You can drill down to see the same information per protection and finally per protected virtual machine.

The Usage tab includes the protections and the virtual machines that are no longer actively used by a tenant but exist on the backup target.

Nodes tab

This tab displays information about NetBackup for OpenStack cluster nodes. The following information is shown:

- Node name
- Node ID
- NetBackup for OpenStack Version of the node
- IP address
- Active controller node (True/False)
- Status of the node

The virtual IP is shown as its own node. It is displayed below the current active controller node.

NBOSDM tab (NetBackup for OpenStack data mover service)

This tab displays the information about NetBackup for OpenStack data mover service. The following information is shown:

- Service name
- Compute node the service is running on.
- Service status from an OpenStack perspective (enabled or disabled)
- Version of the service
- General status
- Last time the status was updated.

Audit tab	<p>Audit logs provide the sequence of protection-related activities that users perform such as protection creation, snapshot creation, and so on. The following information is shown:</p> <ul style="list-style-type: none">■ Date and time of the entry■ What task has been done■ Project the task has performed in.■ User that performed the task. <p>The audit log can be searched for strings. For example, only entries done by a specific user.</p> <p>Additionally, the shown time frame can be changed as necessary.</p>
Protection Plan tab	Use the Protection Plan tab to work with the protection plans.
Settings tab	<p>This tab manages all global settings for the cloud. NetBackup for OpenStack has two types of settings:</p> <ul style="list-style-type: none">■ Email settings See “Configuring the email settings” on page 142.■ Job scheduler settings See “Enabling or disabling a job scheduler” on page 144.

Configuring the email settings

These settings are used by NetBackup for OpenStack to send email reports of recovery points and restores to users. The email settings must be configured to provide email notification to OpenStack users.

The following information is required to configure the email settings:

- SMTP server
- SMTP username
- SMTP password
- SMTP port
- SMTP time-out
- Sender email address

A test email can be sent directly from the configuration page.

To work with email settings through the CLI use the following commands:

Configuring the email settings using the command-line

- 1 Set an email setting for the first time or after the deletion.

```
nbosjm setting-create [--description <description>]
                        [--category <category>]
                        [--type <type>]
                        [--is-public {True,False}]
                        [--is-hidden {True,False}]
                        [--metadata <key=value>]
                        <name> <value>
```

- `--description` Optional description (Default=None). Not required for email settings.
- `--category` Optional setting category (Default=None). Not required for email settings.
- `--type` Settings type. Set to `email_settings`.
- `--is-public` Sets if the setting can be seen publicly. Set to `False`.
- `--is-hidden` Sets if the setting should always be hidden. Set to `False`.
- `--metadata` Sets if the setting can be seen publicly. Not required for email settings.
- `<name>` Name of the setting.
- `<value>` Value of the setting.

2 Update the existing email settings.

```
nbosjm setting-update [--description <description>]
                      [--category <category>]
                      [--type <type>]
                      [--is-public {True,False}]
                      [--is-hidden {True,False}]
                      [--metadata <key=value>]
                      <name> <value>
```

- `--description` Optional description (Default=None). Not required for email settings.
- `--category` Optional setting category (Default=None). Not required for email settings.
- `--type` Settings type. Set to `email_settings`.
- `--is-public` Sets if the setting can be seen publicly. Set to `False`.
- `--is-hidden` Sets if the setting will always be hidden. Set to `False`.
- `--metadata` Sets if the setting can be seen publicly. Not required for email settings.

- `<name>` Name of the setting.
- `<value>` Value of the setting.

3 Show the existing email settings.

```
nbosjm setting-show [--get_hidden {True,False}] <setting_name>
```

- `--get_hidden` Hidden settings (True) or not (False). Not required for email settings, use `False` if set.
- `<setting_name>` Name of the setting to show.

4 Delete the email settings.

```
nbosjm setting-delete <setting_name>
```

- `<setting_name>` Name of the setting to delete.

Table 7-1 Email settings

Setting name	Value type	Example
smtp_default_recipient	String	admin@example.net
smtp_default_sender	String	admin@example.net
smtp_port	Integer	587
smtp_server_name	String	Mailserver_A
smtp_server_username	String	admin
smtp_server_password	String	password
smtp_timeout	Integer	10
smtp_email_enable	Boolean	True

Enabling or disabling a job scheduler

The global job scheduler can be used to deactivate all scheduled policies without modifying each one of them.

To enable or disable a job scheduler using Horizon

- 1 Log in to the Horizon console with the administrator user.
- 2 Navigate to **Admin > NBOS Backup Admin > NetBackupOpenStack > Settings**.
- 3 Click **Disable/Enable Job Scheduler**.

- 4 Select or clear the **Job Scheduler Enabled** box.
- 5 Click **Change** to confirm.

To enable or disable a job scheduler using a command-line

- 1 Get the status of the global job scheduler.

```
nbosjm get-global-job-scheduler
```

- 2 Enable a job scheduler.

```
nbosjm enable-global-job-scheduler
```

- 3 Disable a job scheduler.

```
nbosjm disable-global-job-scheduler
```

Protection plan

NetBackup for OpenStack's tenant driven backup service gives tenants control over backup protections. However, sometimes it may be too much control to tenants and the cloud administrators may want to limit what protections are allowed by tenants. For example, a tenant may exceed its quota by performing full backups at a very high frequency. If every tenant was to pursue such a backup protection, it may affect the resource limits set on the cloud infrastructure. Instead, if the NetBackup administrator can define predefined protection plans and each tenant is only limited to those protections then NetBackup administrators can exert better control over backup service.

List the available protection plans

Using Horizon

To see all available policies in Horizon follow these steps:

1. Log in to the Horizon console with the administrator user.
2. Navigate to **Admin > NBOS Backup Admin > NetBackupOpenStack > Protection Plan**.

The following information is shown for each available protection plan:

- Creation time
- Name
- Description
- Status
- Interval

- Snapshot and backup options
- Keep snapshot for
- Keep backup for
- Action

Using CLI

```
nbosjm protection-plan-list
```

Subscribe a project to a protection plan

Using Horizon

To subscribe a project to a protection plan, perform the following steps.

1. Log in to the Horizon console with the administrator user.
2. Navigate to **Admin > NBOS Backup Admin > NetBackupOpenStack > Protection Plan**.
3. Identify the protection plan to subscribe or unsubscribe a project to.
4. Click **Subscribe/Unsubscribe Projects**.
5. Choose projects to add or remove by using the plus or minus options.
6. Click **Apply**.

Using CLI

```
nbosjm protection-plan-assign [--add_project <project_id>]
                             [--remove_project <project_id>]
                             <protection_id>
```

- `--add_project` ID of the project to assign protection plan to.
- `--remove_project` ID of the project to remove a protection plan from.
- `<protection_id>` Protection plan to be subscribed or unsubscribed.

Managing the trusts

NetBackup for OpenStack uses the OpenStack keystone trust system, which enables the NetBackup for OpenStack service user to act in the name of another OpenStack user.

This system is used during all backup and restore features.

OpenStack administrators should never have the need to directly work with the trusts created. The cloud-trust is created during the NetBackup for OpenStack configuration and further trusts are created as necessary upon creating or modifying a protection.

You can manage the trusts using the command line only.

To manage the trusts

1 List all trusts.

```
nbosjm trust-list
```

2 Show a trust.

```
nbosjm trust-show <trust_id>
```

3 Create a trust.

```
nbosjm trust-create [--is_cloud_trust {True,False}] <role_name>
```

- **<role_name>** Name of the role that trust is created for.
- **--is_cloud_trust** Set to **True** to create cloud admin trust. While creating cloud trust use the same user and tenant which was used to configure NetBackup for OpenStack and keep the role admin.

4 Delete a trust.

```
nbosjm trust-delete <trust_id>
```

- **<trust_id>** ID of the trust to be deleted.

Policy import and migration

Each NetBackup for OpenStack policy has a dedicated owner. The ownership of a policy is defined by:

- **OpenStack User:** The OpenStack User-ID assigned to a policy.
- **OpenStack Project:** The OpenStack Project-ID is assigned to a policy.
- **OpenStack Cloud:** The NetBackup for OpenStack Serviceuser-ID assigned to a policy.

OpenStack users can update the user ownership of a policy by modifying the policy. This ownership ensures that only the owners of a policy are able to work with it.

OpenStack Administrators can reassign policies or reimport policies from older NetBackup for OpenStack installations.

Note: The `policy-reassign` command is not supported in NetBackup for OpenStack 10.4.

Importing the policies

You can import policies from the backup target to the NetBackup for OpenStack database.

The policy import feature is designed to import the policies that are owned by the cloud. It does not import or list any policies that are owned by a different cloud.

To import the policies

- 1 Get a list of policies that can be imported.

```
nbosjm policy-get-importpolicies-list [--project_id <project_id>]
[--storage_type <storage_type>] [--backup_path <backup_path>]
```

- `--project_id` List the policies that belong to the specified project only.
- `--storage_type` The storage type (S3 or NFS), where the policies are stored.
- `--backup_path` The backup storage path where backups are stored.

For S3, `storage_type` and `backup_path` are optional parameters.

- 2 Import the policies into the NetBackup for OpenStack database.

```
nbosjm policy-importpolicy [--policies <policyid>] [--storage_type
<storage_type>] [--backup_path <backup_path>]
```

- `--policyids` Specify policy IDs to import. Repeat option for multiple policies.
- `--storage_type` The storage type (S3 or NFS), where the policies are stored.
- `--backup_path` The backup storage path where backups are stored.

For S3, `storage_type` and `backup_path` are optional parameters.

- 3 Verify that the policies are imported properly.

```
./nbosjm ./nbosjm policy-verify-importedpolicies
```

- `--policyids` Specify policy IDs to verify that the policies are imported properly.
- `--storage_type` The storage type (S3 or NFS), where the policies are stored.
- `--backup_path` The backup storage path where backups are stored.

Before you run import policy commands for S3 storage type, perform the following.

1. Add the following details in the `/etc/nbos/nbosjm.conf` file.

```
vault_s3_auth_version = DEFAULT
vault_s3_access_key_id = << s3_access_key >>
vault_s3_secret_access_key = <<s3_secret_access_key>>
vault_s3_region_name = << s3_region_name >>
vault_s3_bucket = << vault_s3_bucket >>
vault_s3_endpoint_url = << vault_s3_endpoint_url >>
vault_s3_signature_version = default
vault_s3_ssl = False
vault_s3_ssl_cert =
vault_enable_threadpool = True
vault_s3_support_empty_dir = False
[s3fuse_sys_admin]
helper_command = sudo /home/stack/myansible/bin/nbosjm-rootwrap
/etc/nbosjm/rootwrap.conf privsep-helper
```

2. Start the `nbos-object-store` service.

```
systemctl start nbos-object-store
```

3. Check the status of the `nbos-object-store` service. It must be in the running state.

```
systemctl status nbos-object-store
```

Orphaned policies

The definition of an orphaned policy is from the perspective of a specific NetBackup for OpenStack installation. Any policy that is located on the backup target storage, but not known to the NetBackup for OpenStack installation is considered orphaned.

Further is to divide between policies that were previously owned by projects/users in the same cloud or are migrated from a different cloud.

The following CLI command provides the list of orphaned policies:

```
nbosjm policy-get-orphaned-policies-list [--migrate_cloud
{True,False}]
[--generate_yaml {True,False}]
```

- `--migrate_cloud` Set to `True` if you want to list policies from other clouds as well. Default is `False`.

- `--generate_yaml` Set to `True` if you want to generate output file in the YAML file format, which is further used as the input for policy reassign API.

Running this command against a backup target with many policies can take a bit of time. NetBackup for OpenStack reads the complete storage and verifies every found policy against the policies known in the database.

Note: The `policy-get-orphaned-policies-list` command is not supported in NetBackup for OpenStack 10.4.

Disaster recovery

This chapter includes the following topics:

- [About disaster recovery in NetBackup for OpenStack](#)

About disaster recovery in NetBackup for OpenStack

In case of a disaster, perform the following steps for recovery:

To perform the disaster recovery for the different OpenStack cloud

- 1 Perform the disaster recovery on NetBackup.

See the *Disaster recovery* chapter of the *NetBackup Troubleshooting Guide*.

- 2 Reconfigure the NetBackup for OpenStack cluster.

See [“Reconfigure the NetBackup for OpenStack cluster”](#) on page 89.

- 3 Import the protections from NetBackup for OpenStack VM.

```
nbosjm protection-import-to-new-cloud
```

The protections are listed as orphaned protections. The project and user of the OpenStack cloud A do not exist on the OpenStack cloud B.

- 4 List the orphaned protections.

Orphaned protections are the protections that are no longer linked to an active tenant or user within the cloud. To identify and list all orphaned protections that do not have `tenant_id` or `user_id` associated with the current cloud environment, run the following command:

```
nbosjm protection-get-orphaned-protections-list [--migrate_cloud  
{True,False}]
```

- `--migrate_cloud` Set to `True` if you want to list policies from other clouds as well. The default value is `False`.

5 Assign protections to a new tenant or user.

```
nbosjm protection-reassign-protections [--old_tenant_ids
<old_tenant_id>]
                                     [--new_tenant_id
<new_tenant_id>]
                                     --protection_plan_id
<protection_plan_id>
                                     [--user_id <user_id>]
                                     [--migrate_cloud
{True,False}]
                                     [--map_file <map_file>]
```

- `--old_tenant_ids` The IDs of the old tenants from which you want to assign the protections.
- `--new_tenant_id` The ID of the new tenant to which you want to assign the protections.
- `--protection_plan_id` The ID of the protection plan to which you want to assign the protections.
- `--user_id` The user ID to which you want to assign the protections.
- `--migrate_cloud` Set to `True` if you want to assign the protections from other clouds also. The default value is `False`
- `--map_file` The file path with the file name of map file. The file format is YAML.

To perform the disaster recovery for the same OpenStack cloud

1 Perform the disaster recovery on NetBackup.

See the *Disaster recovery* chapter of the *NetBackup Troubleshooting Guide*.

2 Reconfigure the NetBackup for OpenStack cluster.

See [“Reconfigure the NetBackup for OpenStack cluster”](#) on page 89.

3 Import protections from the NetBackup for OpenStack VM.

```
nbosjm protection-import
```


- 4 Run the following command to see the status of the import job.

```
nbosjm get-protection-import-status
```

- 5 After the protection import operation is complete, enable the global job scheduler.

```
nbosjm enable-global-job-scheduler
```

Troubleshooting

This chapter includes the following topics:

- [General Troubleshooting Tips](#)
- [Using the nbosjm CLI tool on the NetBackup for OpenStack Appliance](#)
- [Health check of NetBackup for OpenStack](#)
- [Important log files](#)
- [Troubleshooting NBOSDM container in offline state due to unavailable mount point](#)
- [After restore of the Windows instance, the disk is in an offline state](#)
- [Selective restore from snapshot copy fails](#)
- [A backup fails due to an old nova ID in the universal share path](#)
- [Using the NetBackup support utility in NetBackup for OpenStack](#)
- [Cannot create volumes if the metadata size for physical volume and volume group is small](#)
- [NBOSVM configuration fails if DNS server cannot resolve IP address or IP address is wrong](#)
- [Error when storage unit is created with multiple storage servers](#)
- [Snapshot job fails if the OpenStack image is not accessible to the OpenStack user](#)
- [One-click restore fails if the subnet attached to the instance is not accessible to the OpenStack user](#)
- [The NBOSVM configurator UI does not detect the primary server](#)

- A recovery point name is updated to a default name
- NBOS Backups and NBOS Backup Admin tabs disappear from Horizon UI after stack is updated
- The protection creation fails on the Horizon UI
- The NetBackup for OpenStack services do not start after NBOSVM is restarted
- The NBOSVM is not able to communicate with the nbosdmap on the controller node
- Troubleshooting the OpenStack Keystone authentication failure

General Troubleshooting Tips

Troubleshooting inside a complex environment like OpenStack can be very time-consuming. The following tips help to speed up the troubleshooting process to identify root causes.

What is happening where

OpenStack and NetBackup for OpenStack are divided into multiple services. Each service has a very specific purpose that is called during a backup or recovery procedure. Knowing the function of the service helps to understand where the error is, allowing more focused troubleshooting.

NetBackup for OpenStack cluster

The NetBackup for OpenStack Cluster is the Controller of NetBackup for OpenStack. It receives all protection-related requests from the users.

Every task of a backup or restore process is triggered and managed from here. This includes the creation of the directory structure and initial metadata files on the Backup Target.

During a backup process

During a backup process, the NetBackup for OpenStack cluster is also responsible to gathering the metadata about the backed-up virtual machines and networks from the OpenStack environment. It sends API calls towards the OpenStack endpoints on the configured endpoint type to fetch this information. Once the metadata has been received the NetBackup for OpenStack Cluster writes it as JSON files on the Backup Target.

The NetBackup for OpenStack cluster also sends the Cinder Snapshot command.

During a restore process

During the restore process the NetBackup for OpenStack cluster reads the virtual machine metadata from its Database and uses the metadata to create the Shell for the restore. It sends API calls to the OpenStack environment to create the necessary resources.

nbosdmapi

The nbosdmapi service is the connector between the NetBackup for OpenStack cluster and the data mover running on the compute nodes.

The purpose of the nbosdmapi service is to identify which compute node is responsible for the current backup or restore task. To do so, the nbosdmapi service connects to the nova API database requesting the compute host of a provided virtual machine.

Once the compute host has been identified the nbosdmapi forwards the command from the NetBackup for OpenStack Cluster to the data mover running on the identified compute host.

nbosdm

The nbosdm is the NetBackup for OpenStack service running on the compute nodes.

Each data mover is responsible for the virtual machines running on top of its compute node. A data mover cannot work with virtual machines running on a different compute node.

The data mover controls the freeze and thaw of virtual machines as well as the actual movement of the data.

Everything on the Backup Target happens as the user nova

NetBackup for OpenStack reads and writes on the Backup Target as nova:nova.

The POSIX user-id and group-id of nova:nova need to be aligned between the NetBackup for OpenStack Cluster and all compute nodes. Otherwise backups or restores may fail with permission or file not found issues.

Alternative ways to achieve the goal are possible, as long as all required nodes can fully write, and read as nova:nova on the Backup Target.

It is recommended to verify the required permissions on the Backup Target in case of any errors during the data transfer phase or in case of any file permission errors.

NetBackup for OpenStack Trustee Role

NetBackup for OpenStack uses RBAC to allow the usage of NetBackup for OpenStack features to users.

This trustee role is required and cannot be overwritten using the admin role.

It is recommended to verify the assignment of the NetBackup for OpenStack Trustee Role in case of any permission errors from NetBackup for OpenStack during creation of protections, backups, or restores.

OpenStack Quotas

To protect the Cinder volumes, NetBackup for OpenStack creates Cinder snapshots and additional temporary Cinder volumes. The tenant administrator must configure the OpenStack quotas accordingly to provision adequate snapshots and the volumes that full and incremental backups need. The temporary volumes are used to generate disk map information per disk, and to calculate incrementally changed data.

Volume quota requirement is based on the total number of disks getting backed up simultaneously through one or more protections. As the number of simultaneous backups increases, more volume quota is required. Tenant administrator can determine the volume quota by calculating the sum of the total number of instances and the total number of disks that are attached to those instances. For example, you want to protect 10 instances and each instance has two disks attached. To protect these instances simultaneously through one or more protections, the required volume quota is 30.

Ephemeral disk backup

Ephemeral storage is a non-persistent storage form that is associated only to a specific compute instance. An ephemeral disk that is allocated to an instance gets deleted when the instance is terminated. Ephemeral disks are ideally used to store the temporary data.

NetBackup for OpenStack does not protect the ephemeral disk that is allocated to the virtual machine instance.

Using the nbosjm CLI tool on the NetBackup for OpenStack Appliance

To use the nbosjm CLI tool on the NetBackup for OpenStack appliance it is only necessary to activate the virtual environment of the nbosjm.

```
source /home/stack/myansible/bin/activate
```

An rc-file to authenticate against OpenStack is required.

Health check of NetBackup for OpenStack

NetBackup for OpenStack is composed of multiple services, which can be checked in case of any errors.

On the NetBackup for OpenStack Cluster

nbosjm-policies

This service runs and is active on every NetBackup for OpenStack node.

```
[root@Upstream ~]# systemctl status nbosjm-policies
● nbosjm-policies.service - nbosjm policies service
   Loaded: loaded (/etc/systemd/system/nbosjm-policies.service; enabled;
   vendor preset: disabled)
   Active: active (running) since Wed 2020-06-10 13:42:42 UTC; 1 weeks
   4 days ago
   Main PID: 12779 (nbosjm-wor)
   Tasks: 17
   CGroup: /system.slice/nbosjm-policies.service
           └─12779 /home/stack/myansible/bin/python
/home/stack/myansible/bin/nbosjm-policies
--config-file=/etc/nbosjm/nbosjm.conf
           └─12982 /home/stack/myansible/bin/python
/home/stack/myansible/bin/nbosjm-policies
--config-file=/etc/nbosjm/nbosjm.conf
           └─12983 /home/stack/myansible/bin/python
/home/stack/myansible/bin/nbosjm-policies
--config-file=/etc/nbosjm/nbosjm.conf
           └─12984 /home/stack/myansible/bin/python
/home/stack/myansible/bin/nbosjm-policies
--config-file=/etc/nbosjm/nbosjm.conf
   [...]

```

nbosjm-api

This service runs and is active on every NetBackup for OpenStack node.

```
[root@Upstream ~]# systemctl status nbosjm-api
● nbosjm-api.service - nbosjm api service
   Loaded: loaded (/etc/systemd/system/nbosjm-api.service; disabled;
   vendor preset: disabled)
   Drop-In: /run/systemd/system/nbosjm-api.service.d
            └─50-pacemaker.conf
   Active: active (running) since Thu 2020-04-16 22:30:11 UTC;
   2 months 5 days ago
   Main PID: 11815 (nbosjm-api)
     Tasks: 1
    CGroup: /system.slice/nbosjm-api.service
            └─11815 /home/stack/myansible/bin/python /home/stack/
               myansible/bin/nbosjm-api --config-file=/etc/
               nbosjm/nbosjm.conf
```

nbosjm-scheduler

This service runs and is active on every NetBackup for OpenStack node.

```
[root@Upstream ~]# systemctl status nbosjm-scheduler
● nbosjm-scheduler.service - nbosjm scheduler service
   Loaded: loaded (/etc/systemd/system/nbosjm-scheduler.service; disabled;
   vendor preset: disabled)
   Drop-In: /run/systemd/system/nbosjm-scheduler.service.d
            └─50-pacemaker.conf
   Active: active (running) since Thu 2020-04-02 13:49:22 UTC; 2 months
   20 days ago
   Main PID: 29439 (nbosjm-sch)
     Tasks: 1
    CGroup: /system.slice/nbosjm-scheduler.service
            └─29439 /home/stack/myansible/bin/python /home/stack/myansible
               /bin/nbosjm-scheduler --config-file=/etc/nbosjm/
               nbosjm.conf
```

nbosjm-cron

This service is controlled by pacemaker and runs only on the master node.

```
[root@Upstream ~]# systemctl status nbosjm-cron
```

```

● nbosjm-cron.service - Cluster Controlled nbosjm-cron
   Loaded: loaded (/etc/systemd/system/nbosjm-cron.service; disabled;
   vendor preset: disabled)
   Drop-In: /run/systemd/system/nbosjm-cron.service.d
            └─50-pacemaker.conf
   Active: active (running) since Wed 2021-01-27 19:59:26 UTC; 6 days ago
   Main PID: 23071 (nbosjm-cro)
   CGroup: /system.slice/nbosjm-cron.service
           └─23071 /home/stack/myansible/bin/python3 /home/stack/
myansible/bin/nbosjm-cron --config-file=/etc/nbosjm
/nbosjm.conf
           └─23248 /home/stack/myansible/bin/python3 /home/stack/
myansible/bin/nbosjm-cron --config-file=/etc/nbosjm/
nbosjm.conf

Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant nbosjm-cron
[23071]: ● nbosjm-cron.service - Cluster Controlled nbosjm-cron
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant nbosjm-cron
[23071]: Loaded: loaded (/etc/systemd/system/nbosjm-cron.service; disabled;
vendor preset: disabled)
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant nbosjm-cron
[23071]: Drop-In: /run/systemd/system/nbosjm-cron.service.d
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant nbosjm-cron
[23071]: └─50-pacemaker.conf
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant nbosjm-cron
[23071]: Active: active (running) since Wed 2021-01-27 19:59:26 UTC;
6 days ago
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant nbosjm-cron
[23071]: Main PID: 23071 (nbosjm-cro)
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant nbosjm-cron
[23071]: CGroup: /system.slice/nbosjm-cron.service
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant nbosjm-cron
[23071]: └─23071 /home/stack/myansible/bin/python3 /home/stack/myansible/
bin/nbosjm-cron --config-file=/etc/nbosjm/nbosjm.conf
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant nbosjm-cron
[23071]: └─23248 /home/stack/myansible/bin/python3 /home/stack/myansible/
bin/nbosjm-cron --config-file=/etc/nbosjm/nbosjm.conf
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant nbosjm-cron
[23071]: └─27145 /usr/bin/systemctl status nbosjm-cron

```


Pacemaker Cluster Status

The pacemaker cluster controls and watches the VIP on the NetBackup for OpenStack Cluster. It also controls on which node the nbosjm-api and nbosjm-scheduler service runs.

```
[root@Upstream ~]# pcs status
Cluster name: NetBackup for OpenStack

WARNINGS:
Corosync and pacemaker node names do not match (IPs used in setup?)

Stack: corosync
Current DC: nbosvm1-ansible-ussuri-ubuntu18-vagrant (version
1.1.23-1.el7_9.1-9acf116022) - chapterition with quorum
Last updated: Wed Feb  3 19:20:02 2021
Last change: Wed Jan 27 20:00:12 2021 by root via crm_resource on
nbosvm1-ansible-ussuri-ubuntu18-vagrant

1 node configured
6 resource instances configured

Online: [ nbosvm1-ansible-ussuri-ubuntu18-vagrant ]

Full list of resources:

    virtual_ip      (ocf::heartbeat:IPaddr2):      Started nbosvm1-ansible-
ussuri-ubuntu18-vagrant
    virtual_ip_public (ocf::heartbeat:IPaddr2):      Started nbosvm1-
ansible-ussuri-ubuntu18-vagrant
    virtual_ip_admin  (ocf::heartbeat:IPaddr2):      Started nbosvm1-
ansible-ussuri-ubuntu18-vagrant
    virtual_ip_internal (ocf::heartbeat:IPaddr2):      Started nbosvm1-
ansible-ussuri-ubuntu18-vagrant
    nbosjm-cron       (systemd:nbosjm-cron):      Started nbosvm1-ansible-
ussuri-ubuntu18-vagrant
    Clone Set: lb_nginx-clone [lb_nginx]
        Started: [ nbosvm1-ansible-ussuri-ubuntu18-vagrant ]

Daemon Status:
    corosync: active/enabled
    pacemaker: active/enabled
    pcsd: active/enabled
```

Mount availability

The NetBackup for OpenStack Cluster needs access to the Backup Target and should have the correct mount at all times.

```
[root@Upstream ~]# df -h
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                  3.8G         0   3.8G   0% /dev
tmpfs                     3.8G      38M   3.8G   1% /dev/shm
tmpfs                     3.8G     427M   3.4G  12% /run
tmpfs                     3.8G         0   3.8G   0% /sys/fs/cgroup
/dev/vda1                 40G      8.8G   32G  22% /
tmpfs                     773M         0   773M   0% /run/user/996
tmpfs                     773M         0   773M   0% /run/user/0
10.10.2.20:/upstream      1008G     704G   254G   74% /var/NetBackupOpenStack-mounts/
                               MTAuMTAuMi4yMDovdXBzdHJlYW0=
10.10.2.20:/upstream2     483G      22G   462G    5% /var/NetBackupOpenStack-mounts/
                               MTAuMTAuMi4yMDovdXBzdHJlYW0y
```

The nbosdmapi service

The nbosdmapi service has its own Keystone endpoints, which should be checked in addition to the actual service status.

```
[root@upstreamcontroller ~(keystone_admin)]# openstack endpoint list |
grep nbosdmapi
| 47918c8df8854ed49c082e398a9572be | RegionOne | nbosdmapi
| datamover      | True      | public    | http://10.10.2.10:8784/v2
| cca52aff6b2a4f47bcc84b34647fba71 | RegionOne | nbosdmapi
| datamover      | True      | internal  | http://10.10.2.10:8784/v2
| e9aa6630bfb74a9bb7562d4161f4e07d | RegionOne | nbosdmapi
| datamover      | True      | admin     | http://10.10.2.10:8784/v2

[root@upstreamcontroller ~(keystone_admin)]# curl http://10.10.2.10:8784/v2
{"error": {"message": "The request you have made requires authentication.",
"code": 401, "title": "Unauthorized"}}
```

```
[root@upstreamcontroller ~(keystone_admin)]# systemctl status
nbosdmapi.service
● nbosdmapi.service - NetBackup for OpenStack datamover API service
   Loaded: loaded (/etc/systemd/system/nbosdmapi.service; enabled;
   vendor preset: disabled)
```

```
Active: active (running) since Sun 2020-04-12 12:31:11 EDT; 2 months
9 days ago
Main PID: 11252 (python)
Tasks: 2
CGroup: /system.slice/nbosdmapl.service
└─11252 /usr/bin/python /usr/bin/nbosdmapl-api
└─11280 /usr/bin/python /usr/bin/nbosdmapl-api
```

The nbosdm service

The nbosdm service is running on each compute node and is integrated as nova compute service.

```
[root@upstreamcontroller ~(keystone_admin)]# openstack compute service list
```

```
[root@upstreamcompute1 ~]# systemctl status nbosdm
```

```
● nbosdm.service - NetBackup for OpenStack datamover service
   Loaded: loaded (/etc/systemd/system/nbosdm.service; enabled; vendor
   preset: disabled)
   Active: active (running) since Wed 2020-06-10 10:07:28 EDT; 1 weeks
   4 days ago
   Main PID: 10384 (python)
   Tasks: 21
   CGroup: /system.slice/nbosdm.service
           └─10384 /usr/bin/python /usr/bin/nbosdm --config-file=/etc/nova/
   nova.conf --config-file=/etc/nbosdm/nbosdm.conf
```

Important log files

On the NetBackup for OpenStack Nodes

The NetBackup for OpenStack Cluster contains multiple log files.

The main log is nbosjm-policies.log, which contains all logs about ongoing and past NetBackup for OpenStack backup and restore tasks. It can be found at:

```
/var/log/nbosjm/nbosjm-policies.log
```

The next important log is the nbosjm-api.log, which contains all logs about API calls received by the NetBackup for OpenStack Cluster. It can be found at:

```
/var/log/nbosjm/nbosjm-api.log
```

The log for the third service is the `nbosjm-scheduler.log`, which contains all logs about the internal job scheduling between NetBackup for OpenStack nodes in the NetBackup for OpenStack Cluster.

```
/var/log/nbosjm/nbosjm-scheduler.log
```

The last service running on the NetBackup for OpenStack Nodes is the `nbosjm-cron` service, which controls the scheduled automated backups.

```
/var/log/nbosjm/nbosjm-cron.log
```

NetBackup for OpenStack data mover service logs on RHOSP

Following are the NetBackup for OpenStack data mover service logs on RHOSP:

- **nbosdmapl log**

The log for the NetBackup for OpenStack data mover API service is located on the nodes, typically controller, where the NetBackup for OpenStack data mover API container is running under:

```
/var/log/containers/nbosdmapl/nbosdmapl.log
```

- **nbosdm log**

The log for the NetBackup for OpenStack data mover service is located on the nodes, typically compute, where the NetBackup for OpenStack data mover container is running under:

```
/var/log/containers/nbosdm/nbosdm.log
```

For VxMS-supported Linux file systems, VxMS logs for the incremental backups are stored at the following location: `/usr/opensv/netbackup/logs/vxms/`

VxMS log level is defined in `/usr/opensv/netbackup/bp.conf` file and it is configured to 2 by default.

```
VXMS_VERBOSE = 2
```

You can configure the log level from 0 to 5. A higher number results in more verbose logs.

Note: VxMS log may take significant disk space when the log verbosity is set to high. Ensure that you clean up the VxMS log files periodically to avoid any disk space-related issues.

Table 9-1 VxMS log levels

Log level	Description
0	No logging

Table 9-1 VxMS log levels (*continued*)

Log level	Description
1	Error logging
2	Level 1 + warning messages
3	Level 2 + informative messages
4	Same as level 3.
5	Highly verbose (includes level 1) + auxiliary evidence files (.MMF, .DUMP, .XML, .RVPMEM)

NetBackup for OpenStack data mover service logs on Ansible OpenStack

Following are the NetBackup for OpenStack data mover service logs on Ansible OpenStack:

- **nbosdmapl log**

The log for the NetBackup for OpenStack data mover API service is located on the nodes, typically controller, where the NetBackup for OpenStack data mover API container is running. Log on to the nbosdmapl container using `lxc-attach` command.

```
lxc-attach -n controller_nbosdmapl_container-all984bf
```

The log file is then located under:

```
/var/log/nbosdmapl/nbosdmapl.log
```

- **nbosdm log**

The log for the NetBackup for OpenStack data mover service is typically located on the compute nodes and the logs can be found here:

```
/var/log/nbosdm/nbosdm.log
```

For VxMS-supported Linux file systems, VxMS logs for the incremental backups are stored at the following location: `/usr/opensv/netbackup/logs/vxms/`

VxMS log level is defined in `/usr/opensv/netbackup/bp.conf` file and it is configured to 2 by default.

```
VXMS_VERBOSE = 2
```

You can configure the log level from 0 to 5. A higher number results in more verbose logs.

Note: VxMS log may take significant disk space when the log verbosity is set to high. Ensure that you clean up the VxMS log files periodically to avoid any disk space-related issues.

Table 9-2 VxMS log levels

Log level	Description
0	No logging
1	Error logging
2	Level 1 + warning messages
3	Level 2 + informative messages
4	Same as level 3.
5	Highly verbose (includes level 1) + auxiliary evidence files (.MMF, .DUMP, .XML, .RVPMEM)

NetBackup for OpenStack data mover service logs on Kolla

- **nbosdmapi log:**
The log for the NetBackup for OpenStack data mover API service is located on the nodes, typically controller, where the NetBackup for OpenStack data mover API container is running.
Log in to the nbosdmapi container using docker command.

```
docker container exec -it < nbosdmapi_container_id > /bin/bash
```


The log file is then located under: `/var/log/kolla/nbosdmapi/nbosdmapi.log`
- **nbosdm log:**
The log for the NetBackup for OpenStack data mover service is typically located on the compute nodes.
Log into the nbosdm container using docker command.

```
docker container exec -it < nbosdm_container_id > /bin/bash
```


The log file is then located under: `/var/log/kolla/nbosdm/nbosdm.log`

Troubleshooting NBOSDM container in offline state due to unavailable mount point

If NetBackup for OpenStack data mover container stops responding, it could be due to the unavailable mount point or incorrect mount path.

Check the logs for an error. NetBackup for OpenStack data mover container logs are stored at the following location:

- RHOSP: `/var/log/nbosdm/nbosdm.log`
- OpenStack Ansible: `/var/log/nbosdm/nbosdm.log`

Example log file:

```
2021-08-31 12:42:37.630 17 ERROR
oslo_messaging.rpc.server nbosdm.exception.InvalidNFSMountPoint:
Error: '/var/lib/nova/NetBackupOpenStack-mounts/MTAuMjIxLjk5LjUx
Oi9tbnQvbmZzX3NoYXJlL2RvY3M=' is not '10.2xx.xx.50:/mnt/nfs_share/docs'
mounted
2021-08-31 12:42:37.630 17 ERROR oslo_messaging.rpc.server
```

To resolve this issue on RHOSP

- 1 Specify the correct mount path in `nbos_env.yaml` file.
- 2 Run the following deployment command:

```
openstack overcloud deploy
```

To resolve this issue on OpenStack Ansible

- 1 Uninstall NBOSDM and NBOSDMAPI service.


```
openstack-ansible os-nbos-install.yml --tags "nbos-all-uninstall"
```
- 2 Specify the correct mount path in the `/etc/openstack_deploy/user_nbos_vars.yml` file.
- 3 Run the following installation command:


```
openstack-ansible os-nbos-install.yml
```

After restore of the Windows instance, the disk is in an offline state

When you restore the Windows instance, the disk that is attached to the instance is in an offline state. The disk does not appear online automatically for Windows instances after the restore.

To make the disk appear online automatically after the restore, update SAN policy to **OnlineAll** before backup of the instance.

To update the SAN policy

- 1 Run Windows command prompt as an administrator.
- 2 Type `diskpart` and press **Enter**.
- 3 Type `san` and press **Enter** to view the current SAN policy.
- 4 Type `san POLICY=OnlineAll` and press **Enter** to update the SAN policy to **OnlineAll**.

Note: This issue is applicable only for Windows Server 2022, Windows Server 2019, Windows Server 2016, Windows Server 2012 R2, and Windows Server 2012.

Selective restore from snapshot copy fails

Selective restore from snapshot copy may fail for nova-booted instances with the following error: `copy_backup_image_to_volume operation failed`.

The selective restore from snapshot copy fails if the compute node or the hypervisor that is selected by OpenStack to create a new instance differs from the compute node where the original instance resides. In this case, perform a selective restore from the backup copy.

A backup fails due to an old nova ID in the universal share path

If the universal share path has the old nova ID, the backup job fails.

A backup job also fails with the following error message on the Horizon UI:

```
Failed taking backup of policy snapshot: 'NoneType' object has no
attribute 'strip'
```

To resolve this issue

- 1 Stop the following services.

```
systemctl stop nbosjm-policies
systemctl stop nbosjm-api
systemctl stop nbosjm-scheduler
systemctl stop nbosjm-cron
```

- 2 Run the script `/home/stack/nova_userid.sh` to change the nova ID.

```
./nova_userid.sh
```


3. Run the following command to change the directory ownership to nova for the directories `/etc/nbosjm` and a mount directory, for example `/var/nbos`.

```
chown -R nova:nova <directory_name>
```

- 4 Restart the following services.

```
systemctl stop nbosjm-policies
systemctl stop nbosjm-api
systemctl stop nbosjm-scheduler
systemctl stop nbosjm-cron
```

Using the NetBackup support utility in NetBackup for OpenStack

The NetBackupsupport utility (nbsu) is a command line tool. It queries the host and gathers appropriate diagnostic information about NetBackup and the operating system.

You can use this utility to gather diagnostic information about the NBOSVM. It collects all the log files that are generated under the `/var/log/` directory on NBOSVM and creates a `.tgz` file. You can use this information to troubleshoot the issues.

Note: The NetBackup support utility must be run on NBOSVM only.

To use the NetBackup support utility

- 1 Log on to the NetBackup for OpenStack virtual machine.
- 2 Change the directory to `/usr/opensv/netbackup/bin/support`.
- 3 Run the utility with NBOSVM role.

```
./nbsu -r nbosvm
```

A `.tgz` file is created, which contains all the logs available in the `/var/log` directory.

For example: `NBSU_<hostname>_nbosvm_10092023_082422.tgz`

Cannot create volumes if the metadata size for physical volume and volume group is small

Volumes cannot be created if the metadata size that is provided for physical volume and the volume group is not sufficient.

To resolve this issue, provide the sufficient metadata size while creating the physical volume and the volume group.

To check the metadata size of the volume, run the following commands:

```
pvdisplay -C -o name,mda_size,mda_free
vgdisplay -C -o name,mda_size,mda_free
```

While creating a physical volume or the volume group, run the following command to set the metadata size:

```
pvcreate -metadatasize <metadata size>
```

For example, `pvcreate --metadatasize 1g`

NBOSVM configuration fails if DNS server cannot resolve IP address or IP address is wrong

NBOSVM configuration fails if the wrong IP address is configured in the `/etc/hosts` file. It also fails if the DNS server cannot resolve the IP address.

On the NBOSVM configurator UI, if the Controller nodes field has the wrong IP or short name, NBOSVM configuration fails.

Ensure that the Controller nodes field has the correct IP or short name for all the NBOSVM nodes.

Error when storage unit is created with multiple storage servers

When you create a storage unit with the multiple storage servers, you may get the network error if NetBackup selects the storage server that does not have universal share configured.

If you create a storage unit with the multiple storage servers, ensure that all the storage servers have all the media servers configured.

To resolve this issue

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Disk storage**.
- 3 Click the **Storage servers** tab.
- 4 Click the storage server.
- 5 Under **Media servers**, add all other media servers that are in the storage unit.
- 6 On the left, click **Hosts > Hosts properties**.
- 7 Select the media server and click **Edit media server**.
- 8 Under **Servers > Additional servers** click **Add** to add all other media servers that are in the storage unit.

Snapshot job fails if the OpenStack image is not accessible to the OpenStack user

If the OpenStack image is not accessible to the OpenStack user who triggers snapshot job, the snapshot job fails with the following error message:

```
Couldn't find the image <image_id> of instance <instance_id> for
snapshot_id <snapshot_id>. Check whether image is present or has
required permission for user <user_id> and project <project_id>
```

To resolve this issue, ensure that the OpenStack image is accessible to the OpenStack user.

One-click restore fails if the subnet attached to the instance is not accessible to the OpenStack user

If the subnet attached to the OpenStack instance is not accessible to the OpenStack user who triggers one-click restore, one-click restore fails with the following error message:

```
Tenant <project_id> not allowed to create port on this network
```

To resolve this issue, ensure that the subnet that is attached to the OpenStack instance is accessible to the OpenStack user.

The NBOSVM configurator UI does not detect the primary server

The NBOSVM configurator UI cannot detect the primary server when master server's FQDN is not added in `/etc/hosts` file on the NetBackup for OpenStack VM.

To resolve this issue, add the primary server name in the `/etc/hosts` file on the NetBackup for OpenStack VM.

A recovery point name is updated to a default name

While creating a recovery point, if you name it with a single character, it takes the default name "recovery point" after the upgrade or the auto image replication.

To resolve this issue, ensure that you use two or more characters for a recovery name.

NBOS Backups and NBOS Backup Admin tabs disappear from Horizon UI after stack is updated

After you update the OpenStack stack, the **NBOS Backups** and **NBOS Backup Admin** tabs disappear from the NetBackup for OpenStack Horizon UI. The stack update incorrectly removes the endpoints from OpenStack.

To resolve this issue, run the script `register_nbopenstack_service.sh` on the director node. This script is provided with the installation packages and is available at the following path: `<download location> /nbos-cfg-scripts/redhat-director scripts/rhosp17.1/register_nbopenstack_service.sh`

```
sh register_nbopenstack_service.sh {overcloudrc file} {NBOS Protocol}
{NBOSVM VIP}
```

For example,

```
sh register_nbopenstack_service.sh /home/stack/overcloudrc http
10.xxx.xxx.xx
```

This script registers the NetBackup for OpenStack service so that the **NBOS Backups** and **NBOS Backup Admin** tabs appear on the Horizon UI.

The protection creation fails on the Horizon UI

While creating a protection, it fails with the following error:

```
Error: subscription request failed with status 404.
```

Verify if the protection plan is deleted on the NetBackup web UI. If the protection plan is deleted, use another protection plan to create the protection.

The NetBackup for OpenStack services do not start after NBOSVM is restarted

The following NetBackup for OpenStack services do not start after NBOSVM is restarted:

- Nginx
- rabbitmq-server

To resolve this issue

- 1 Run the following commands to enable and start the services on all the three NBOSVM nodes.

```
systemctl enable nginx
systemctl enable rabbitmq-server
```

```
systemctl start nginx
systemctl start rabbitmq-server
```

- 2 Run the following commands on one of the NBOSVM nodes to start the NBOSVM cluster services:

```
pcs resource cleanup
pcs resource refresh
```

The NBOSVM is not able to communicate with the nbosdmapl on the controller node

If the NBOS VM port is blocked on the controller node, NBOS VM is not able to communicate with the nbosdmapl on the controller node.

If the NBOS VM is configured with HTTPS, the port is 13784. If NBOSVM is configured with HTTP, the port is 8784.

To enable the port on all the controller nodes

- 1 Run the following command to identify the DROP iptables rule line number:

```
iptables -L --line-numbers | grep -i DROP
```

- 2 Run the following command to insert the iptables rule before the DROP rule.

```
sudo iptables -I INPUT <linenumber> -p tcp -s <nbosvm subnet>
--dport <HTTP/HTTPS port number> -m conntrack --ctstate
NEW,ESTABLISHED -j ACCEPT
```

For example, if the DROP iptables rule line number is 88, the NBOSVM subnet is 10.xxx.xxx.xx/20, and NBOSVM is configured with HTTPS, the command is:

```
sudo iptables -I INPUT 87 -p tcp -s 10.xxx.xxx.xx/20 --dport 13784
-m conntrack --ctstate NEW,ESTABLISHED -j ACCEPT
```

Troubleshooting the OpenStack Keystone authentication failure

The OpenStack Keystone authentication fails with the following error message:

```
The request you have made requires authentication. (HTTP 401)
```

To resolve this issue

- 1 Source the OpenStack RC file for the OpenStack project and the user for which the issue appears.

```
source <OpenStack RC file path>
```

For example, source /home/openrc.sh

- 2 List the trust ID for the OpenStack project and the user for which the issue appears.

```
nbosjm trust-list
```

- 3 Run the following command to delete the trust:

```
nbosjm trust-delete <TrustID>
```