

Veritas NetBackup™ for OpenStack Administrator's Guide

Release 9.1

VERITAS™

Veritas NetBackup™ for OpenStack Administrator's Guide

Legal Notice

Copyright © 2021 Veritas Technologies LLC. All rights reserved.

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

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

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

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

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

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

Veritas Technologies LLC
2625 Augustine Drive.
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within the company to answer your questions in a timely fashion.

Our support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about our support offerings, you can visit our website at the following URL:

www.veritas.com/support

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.veritas.com/support

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information
- Available memory, disk space, and NIC information

- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Technical Support
 - Recent software configuration changes and network changes

Licensing and registration

If your product requires registration or a license key, access our technical support Web page at the following URL:

www.veritas.com/support

Customer service

Customer service information is available at the following URL:

www.veritas.com/support

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Advice about technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs, DVDs, or manuals

Support agreement resources

If you want to contact us regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Contents

Technical Support	4	
Chapter 1	Introduction	12
	About NetBackup for OpenStack	12
	NetBackup for OpenStack Architecture	13
	Backup as a Service	13
	Main Components	14
	Service Endpoints	15
	Network topology	16
Chapter 2	Deploying NetBackup for OpenStack	17
	Requirements	17
	System requirements NetBackup for OpenStack VM	18
	NetBackup for OpenStack network considerations	19
	Existing endpoints in OpenStack	19
	OpenStack endpoints required by NetBackup for OpenStack	19
	Recommendation: Provide access to all OpenStack Endpoint types	20
	Backup target access required by NetBackup for OpenStack	20
	Example of a typical NetBackup for OpenStack network integration	21
	Other examples of NetBackup for OpenStack network integrations	23
	Preparing the installation	25
	Tenant Quotas	25
	AWS S3 eventual consistency	25
	NetBackup for OpenStack Cluster	26
	Spinning up the NetBackup for OpenStack VM	26
	Creating the cloud-init image	26
	Spinning up the NetBackup for OpenStack appliance	28
	Uninstalling cloud-init after first start	28
	Installing NetBackup for OpenStack Components	28
	Installing on RHOSP	29

- Configuring NetBackup for OpenStack 38
 - Details needed for the NetBackup for OpenStack Appliance 39
 - Advanced settings 44
 - Starting the configurator 45
- Post Installation Health-Check 46
 - Verify the NetBackup for OpenStack Appliance services are up
..... 46
 - Check the NetBackup for OpenStack pacemaker and NGINX
cluster 48
 - Verify API connectivity of the NetBackup for OpenStack Appliance
..... 48
 - Verify the nbosdm services are up and running 49
 - Verify that the NFS Volume is correctly mounted 50
- Uninstalling from RHOSP 51
 - Clean NetBackup for OpenStack Datamover API service 52
 - Clean NetBackup for OpenStack Datamover Service 53
 - Clean NetBackup for OpenStack haproxy resources 54
 - Clean NetBackup for OpenStack Keystone resources 55
 - Clean NetBackup for OpenStack database resources 55
 - Revert overcloud deploy command 56
 - Revert back to original RHOSP Horizon container 56
 - Destroy the NetBackup for OpenStack VM Cluster 56
- Install workloadmgr CLI client 57
 - About the workloadmgr CLI client 57
 - Installing the workloadmgr client 57

Chapter 3 Configuring NetBackup OpenStack Appliance 59

- Reconfigure the NetBackup for OpenStack Cluster 59
- Configuring the NetBackup master server details 60
 - About security management and certificates in NetBackup 61
- Change NetBackup for OpenStack dashboard password 62
- Reset NetBackup for OpenStack dashboard password 62
- Reinitialize NetBackup for OpenStack 62
- Download NetBackup for OpenStack logs 62

Chapter 4 Configuring NetBackup Master Server 64

- License for OpenStack plug-in for NetBackup 64
- Allow NetBackup for OpenStack VM on NetBackup master server
..... 64
- About launching the OpenStack Horizon UI from the NetBackup web
UI 65

	Adding the OpenStack Horizon instance on NetBackup web UI	66
	Creating the custom role for NetBackup for OpenStack administrator	66
	Launching the Horizon UI from the NetBackup web UI	67
Chapter 5	NetBackup for OpenStack policies	68
	About policies	68
	List of policies	68
	Create a policy	69
	Policy overview	70
	Edit a policy	72
	Delete a policy	73
	Unlock a policy	74
	Reset a policy	74
Chapter 6	Performing backups and restores of OpenStack	76
	About snapshots	77
	List of snapshots	77
	Creating a snapshot	79
	Snapshot overview	80
	Delete snapshots	81
	Snapshot Cancel	82
	About restores	83
	List of Restores	83
	Restores overview	83
	Delete a Restore	85
	Cancel a Restore	86
	One-Click Restore	87
	Selective Restore	88
	Inplace Restore	90
	Required restore.json for CLI	91
	General required information	92
	Selective Restore required information	93
	Inplace Restore required information	98
	About file search	99
	Navigating to the file search tab in Horizon	99
	Configuring and starting a file search in Horizon	99
	Choose the VM the file search shall run against	100
	Set the file path	100
	Define the Snapshots to search in	100

- Start the File Search and retrieve the results in Horizon 101
- Doing a CLI File Search 101
- About snapshot mount 102
- Create a File Recovery Manager Instance 102
 - Steps to apply on CentOS and RHEL cloud-images 103
- Mounting a snapshot 103
- Accessing the File Recovery Manager 105
- Identifying mounted snapshots 105
- Unmounting a snapshot 106
- About schedulers 107
- Disable a schedule 107
- Enable a schedule 108
- Modify a schedule 108
- About email notifications 108
- Requirements to activate email Notifications 108
- Activate/Deactivate the email Notifications 109

Chapter 7

- Performing Backup Administration tasks 110**
 - NBOS Backup Admin Area 110
 - Access the NBOS Backup Admin area 110
 - Status overview 111
 - Policies tab 111
 - Usage tab 112
 - Nodes tab 112
 - NBOSDM tab (NetBackup for OpenStack Datamover Service) 112
 - Storage tab 113
 - Audit tab 113
 - Policy Attributes tab 114
 - Settings tab 114
 - Policy Attributes 117
 - List the available policies 117
 - Create policy attributes 118
 - Edit the policy attribute 119
 - Assign/Remove a policy 120
 - Delete a policy 121
 - Policy Quotas 121
 - Work with Policy Quotas via Horizon 122
 - Work with Policy Quotas via CLI 122
 - Managing Trusts 124
 - List all trusts 125
 - Show a trust 125

Create a trust	125
Delete a trust	125
Policy import and migration	125
Import policies	126
Orphaned policies	126
Reassigning policies	127
Disaster Recovery	128
Disaster Recovery Process	129
Mount-paths	129
Example runbook for disaster recovery using NFS	130
Scenario	130
Prerequisites for the disaster recovery process	131
Disaster recovery of a single policy	132
Disaster recovery of a complete cloud	139
Chapter 8 Troubleshooting	150
General Troubleshooting Tips	150
What is happening where	150
Everything on the Backup Target happens as user nova	151
NetBackup for OpenStack Trustee Role	152
OpenStack Quotas	152
Using the workloadmgr CLI tool on the NetBackup for OpenStack Appliance	152
Health check of NetBackup for OpenStack	152
On the NetBackup for OpenStack Cluster	153
The nbosdmapi service	157
The datamover service	157
Important log files	158
On the NetBackup for OpenStack Nodes	158
NetBackup for OpenStack Datamover service logs on RHOSP	159
NetBackup for OpenStack Datamover service logs on Ansible OpenStack	159
Index	160

Introduction

This chapter includes the following topics:

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

About NetBackup for OpenStack

Veritas 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 snapshots. These snapshots can be held in a variety of storage environments including NFS AWS S3 compatible storage. 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 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 snapshots 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.

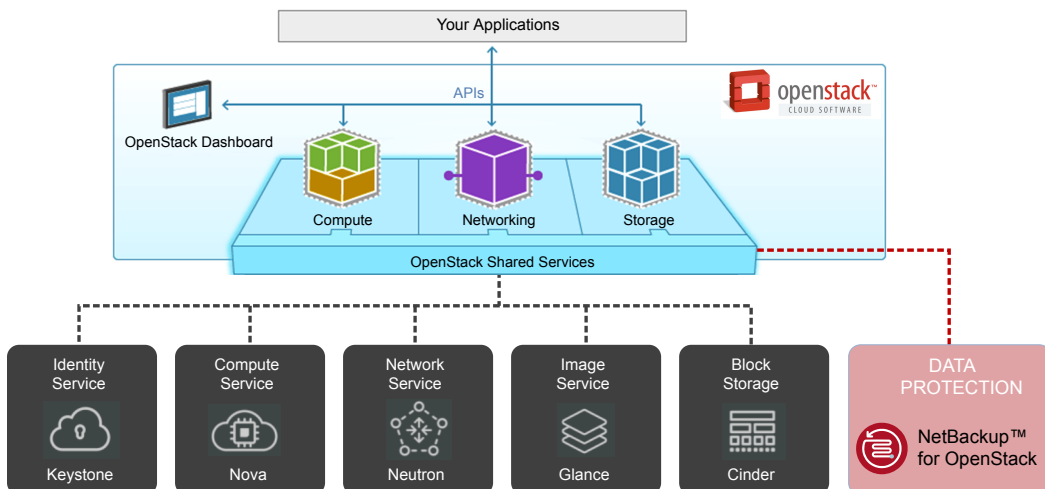
- Faster and reliable recovery. When your applications become complex that snap multiple VMs and storage volumes, our efficient recovery process brings your application from zero to operational with the click of a button.
- Easy migration of workloads between clouds. NetBackup for OpenStack captures all the details of your application and hence our migration includes your entire application stack without leaving anything for guess work.
- 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	Backup as a Service
Main components	Main Components
Service endpoints	Service Endpoints
Network topology	Network topology

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 workloads. NetBackup for OpenStack is very similar to other OpenStack services including Nova, Cinder, Glance, and adheres to all tenets of OpenStack. It is a stateless service that scales with your cloud.

Main Components

Figure 1-2 NetBackup for OpenStack architecture overview

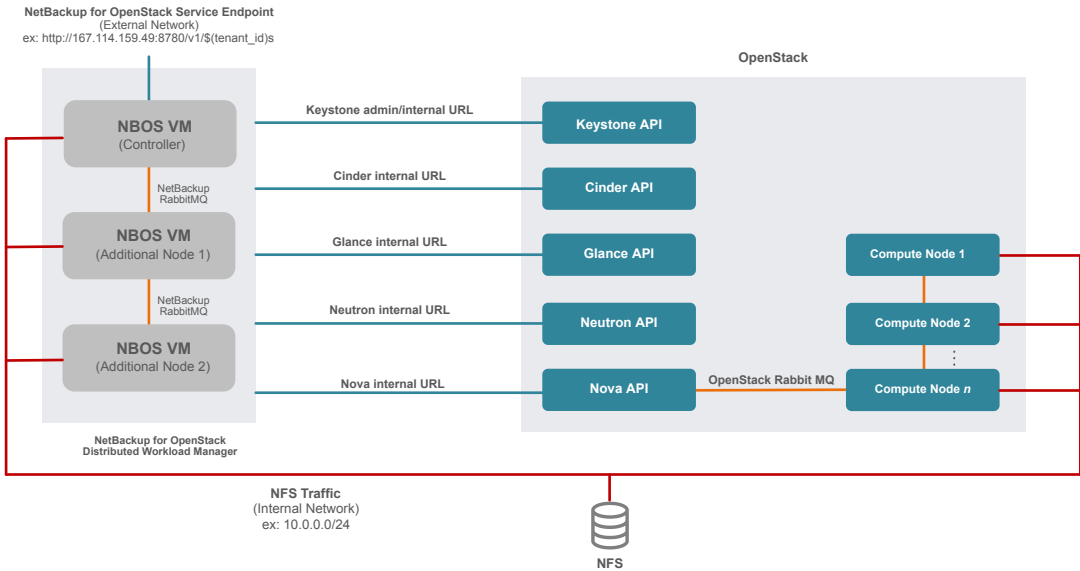


NetBackup for OpenStack has four main software components:

1. NetBackup for OpenStack ships as a QCOW2 image. User can instantiate one or more VMs 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

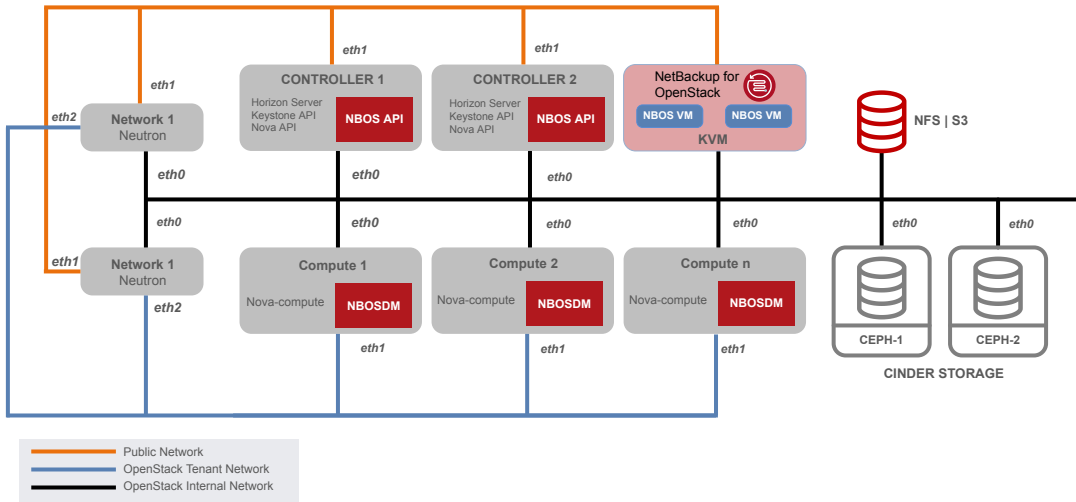
Figure 1-3 Service endpoints overview



NetBackup for OpenStack is both a provider and consumer into 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 and data movers typically use either the internal network or dedicated backup network for storing and retrieving backup images from the backup store.

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 VM](#)
- [Installing NetBackup for OpenStack Components](#)
- [Configuring NetBackup for OpenStack](#)
- [Post Installation Health-Check](#)
- [Uninstalling from RHOSP](#)
- [Install workloadmgr CLI client](#)

Requirements

NetBackup for OpenStack has four main software components:

1. NetBackup for OpenStack ships as a QCOW2 image. User can instantiate one or more VMs from the QCOW2 image on standalone KVM boxes.
2. NetBackup for OpenStack 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 NetBackup for OpenStack VM](#)” on page 18.

See “[Software Requirements](#)” on page 18.

System requirements NetBackup for OpenStack VM

The NetBackup for OpenStack VM gets delivered as a qcow2 image, which gets attached to a virtual machine.

Veritas supports only KVM-based hypervisors.

Note: The NetBackup for OpenStack VM is not supported as instance inside NetBackup for OpenStack.

The recommended size of the VM for the NetBackup for OpenStack Appliance is:

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

vCPU	4
------	---

RAM	24 GB
-----	-------

The qcow2 image itself defines the 40GB disk size of the VM.

In the rare case of the NetBackup for OpenStack VM database or log files getting larger than 40GB disk, contact or open a ticket with Veritas customer support to attach another drive to the NetBackup for OpenStack VM.

Software Requirements

NetBackup for OpenStack has been tested and verified

Software	Version
-----------------	----------------

CentOS	7.9
--------	-----

Virsh	libvirt 2.0.0 and later
-------	-------------------------

QEMU	2.0.0 and later
------	-----------------

Qemu-img	2.6.0 and later
----------	-----------------

Additionally, it is necessary for NFS backup targets to have the `nfs-common` packages installed on the compute nodes.

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 meant to be used by the OpenStack users to work with OpenStack. Internal endpoints are meant to be used by the OpenStack services to communicate with each other. Admin endpoints are meant to be used by OpenStack administrators.

Out of those 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 please 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. Which endpoint type NetBackup for OpenStack uses to communicate with OpenStack is decided during the configuration of the NetBackup for OpenStack appliance.

An exception: The NetBackup for OpenStack Appliance always requires access to the Keystone admin endpoint.

The following network requirement 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

Veritas 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 VM.

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 Datamover 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 VMs 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 Veritas' 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



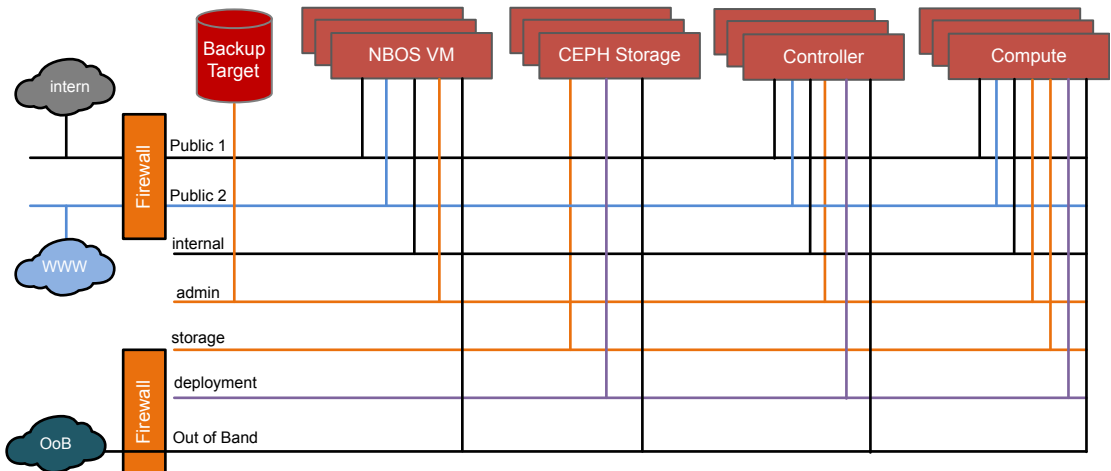
You can combine networks as necessary. As long as the required network access is available NetBackup for OpenStack works.

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:

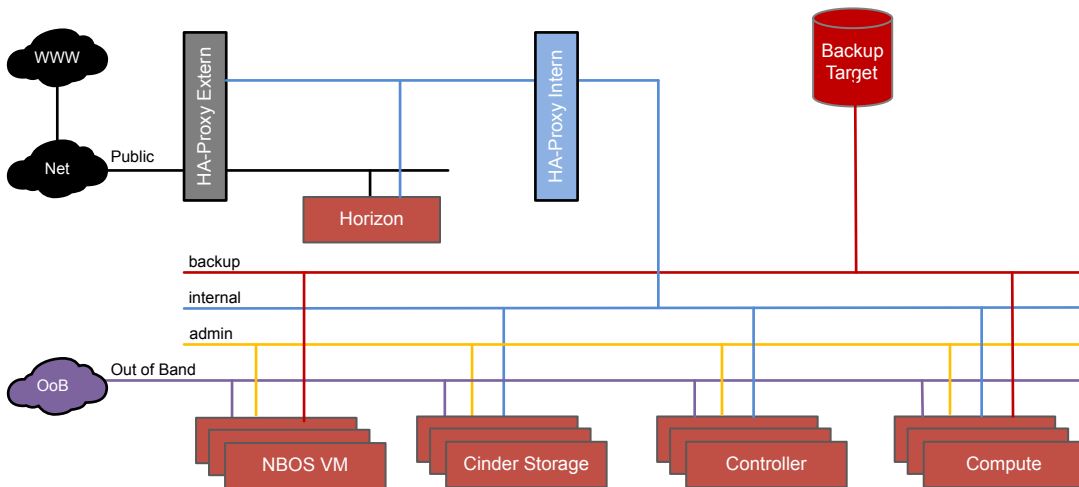
The first example is from a manufacturing company, which wanted to split the networks by function and decided to put the NetBackup for OpenStack backup target on the internal network as the backup and recovery function was identified as an OpenStack internal solution. This example looks complex but integrates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

Figure 2-3 The split them all network example



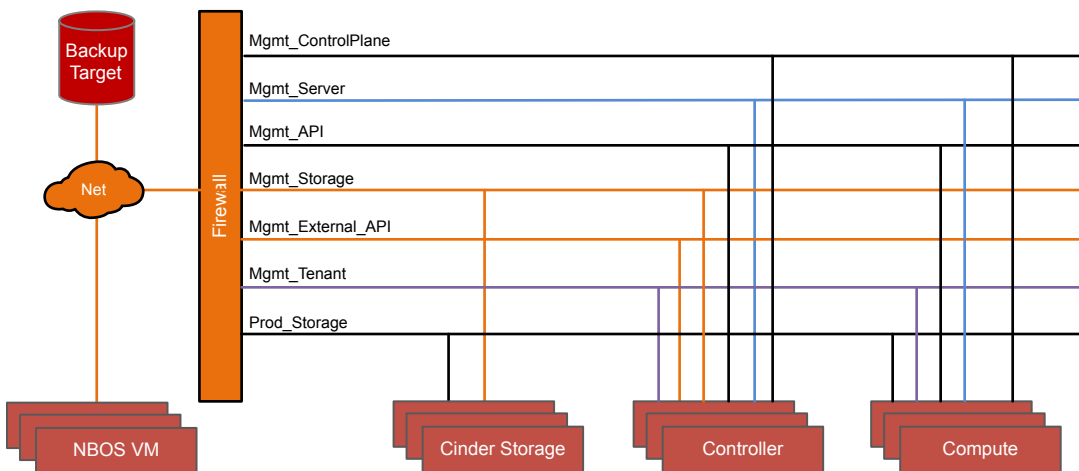
The second example is from a financial institute that wanted to be sure that the OpenStack Users have no direct uncontrolled network access to the OpenStack infrastructure. Following this example requires additional work as the internal HA-Proxy needs to be configured to correctly translate the API calls 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 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 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 sufficient quota for Nova instances and Cinder volumes. Otherwise restore operations result in failures.

AWS S3 eventual consistency

AWS S3 object consistency model includes:

1. Read-after-write
2. Read-after-update
3. Read-after-delete

Each of them describes how an object reaches its consistent state after an object is created, updated, or deleted. None of them provides strong consistency and there is a lag time for an object to reach the consistent state. Though NetBackup for OpenStack employed mechanisms to work around the limitations of eventual consistency of AWS S3, when an object reaches its consistency state is not deterministic. There is no official statement from AWS on how long it takes for an object to reach consistent state. However read-after-write has a shorter time to reach the consistency compared to other IO patterns. Our solution is designed to maximize read-after-write IO pattern. The time in which an object reaches eventual consistency also depends on the AWS region. For example, aws-standard region

does not have strong consistency model compared to us-east or us-west. We suggest using these regions when you create s3 buckets for NetBackup for OpenStack. Though read-after-update IO pattern is hard to avoid completely, we employed ample delays in accessing objects to accommodate larger durations for objects to get into consistent state. However in rare occasions, backups may still fail and need to be restarted.

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 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 VM

The NetBackup for OpenStack Appliance is delivered as qcow2 image and runs as VM 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 either 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 and centos
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. 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 VM name in `virsh`.

The second file is called `user-data` and contains little scripts and information to set up for example the user passwords. 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 has been created the NetBackup for OpenStack appliance can be spun up on the desired KVM server.

Following example command shows how to spin up the NetBackup for OpenStack appliance using virsh 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
```

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
```

Installing NetBackup for OpenStack Components

Once the NetBackup for OpenStack VM or the Cluster of NetBackup for OpenStack VMs has been spun, the actual installation process can begin. This process contains the following steps:

1. Install the NetBackup for OpenStack nbosdmapi service on the control plane.
2. Install the NetBackup for OpenStack 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.

1. Prepare for deployment

1.1] Select "backup target" type

Backup target storage is used to store the backup images that are taken by NetBackup for OpenStack and the details that are needed for configuration:

NetBackup for OpenStack supports the following backup target types

- NFS
 - NFS server is configured
 - NFS share path
- Universal Share
 - Universal Share server is configured
 - NFS Share path

For more information about Universal Share, see *Veritas NetBackup™ Administrator's Guide, Volume I*.
- Amazon S3
 - S3 Access Key
 - Secret Key
 - Region
 - Bucket name
- Other S3 compatible storage (Like Ceph based S3)
 - S3 Access Key
 - Secret Key

- Region
- Endpoint URL (Valid for S3 other than Amazon S3)
- Bucket name

1.2] Copy nbos-cfg-scripts to the undercloud

The following steps are to be done on "undercloud" node on an already installed RHOSP environment. The overcloud-deploy command has to be run successfully already and the overcloud should 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:

rhosp16 rhosp16.1

1.3] If backup target type is "Ceph based S3" with SSL:

If your backup target is ceph S3 with SSL and SSL certificates are self-signed or authorized by private CA, then user needs to provide CA chain certificate to validate the SSL requests. User needs to rename his CA chain cert file to "s3-cert.pem" and copy it into directory - 'nbos-cfg-scripts/redhat-director-scripts/redhat-director-scripts/<RHOSP_RELEASE__Directory/puppet/nbos/files'

```
cp s3-cert.pem /home/stack/nbos-cfg-scripts/
redhat-director-scripts/<RHOSP_RELEASE_DIRECTORY>/puppet/nbos/files/
```

2] Upload 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
```

3] Update 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".

3.1] Add 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::nbosdmapi'
```

3.2] Add 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'
```

4] Prepare 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-rhel8/openstack-keystone:16.0-82
| |
docker://nbos-undercloud:8787/rhosp-rhel8/openstack-barbican-keystone-listener:16.0-84
```

'CONTAINER_TAG' format for RHOSP16: <NBOS_VERSION>-rhosp16

'CONTAINER_TAG' format for RHOSP16.1: <NBOS_VERSION>-rhosp16.1

Example,

```
sudo ./prepare_nbos_images.sh nbos-undercloud 9.0.1017-rhosp16.1
/home/stack/nbos/nbos-rhosp16.1-9.0.1017
```

The changes can be verified using the following commands.

```
(undercloud) [stack@nbos-undercloud scripts]$ sudo podman images |
grep 9.0.1017-rhosp16.1
localhost/nbos-horizon-plugin 9.0.1017-rhosp16.1 8705f72da6d4
5 days ago 1.16 GB
localhost/nbosdmapi 9.0.1017-rhosp16.1 2da0be5dcacb
5 days ago 1.46 GB
localhost/nbosdm 9.0.1017-rhosp16.1 d6e1168faae2
5 days ago 2.97 GB
```

```
(undercloud) [stack@host scripts]$ grep 'Image'
```



```
../environments/nbos_env.yaml
  DockerNbosdmImage: nbos-undercloud:8787/nbosdm:9.0.1017-rhosp16.1
  DockerNbosdmapiImage: nbos-undercloud:8787/nbosdmapi:9.0.1017-rhosp16.1
  ContainerHorizonImage: nbos-undercloud:8787/nbos-horizon-plugin:
9.0.1017-rhosp16.1
```

5] Provide environment details in nbos_env.yaml

Provide backup target details and other 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 datamover API
- datamover password
- Backup target type {nfs/s3}
- In case of NFS
 - List of NFS Shares
 - NFS options
- In case of S3
 - S3 type {amazon_s3/ceph_s3}
 - S3 Access key
 - S3 Secret key
 - S3 Region name
 - S3 Bucket
 - S3 Endpoint URL
 - S3 Signature Version
 - S3 Auth Version
 - S3 SSL Enabled {true/false}
 - S3 SSL Cert

Note: Use ceph_s3 for any non-aws S3 backup targets.

```
resource_registry:
  OS::TripleO::Services::nbosdm: ../services/nbosdm.yaml
  OS::TripleO::Services::nbosdmap: ../services/nbosdmap.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:
    nbosdmapNetwork: internal_api

  ## NetBackup for OpenStack Datamover Password for keystone and database
  nbosdmPassword: "test1234"

  ## NetBackup for OpenStack container pull urls
  DockerNbosdmImage: nbos-undercloud:8787/nbosdm:9.0.1017-rhosp16.1
  DockerNbosdmapImage: nbos-undercloud:8787/nbosdmap:9.0.1017-rhosp16.1

  ## 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:
  9.0.1017-rhosp16.1

  ## Backup target type nfs/s3, used to store snapshots taken by
  NetBackup for OpenStack
  BackupTargetType: 'nfs'

  ## For backup target 'nfs'
  NfsShares: '192.168.122.101:/opt/nbos'
  NfsOptions: 'nolock,soft,timeo=180,intr,lookupcache=none'

  ## For backup target 's3'
  ## S3 type: amazon_s3/ceph_s3
  S3Type: 'amazon_s3'
```

```
## S3 access key
S3AccessKey: ''

## S3 secret key
S3SecretKey: ''

## S3 region, if your s3 does not have any region, just keep the
parameter as it is
S3RegionName: ''

## S3 bucket name
S3Bucket: ''

## S3 endpoint url, not required for Amazon S3, keep it as it is
S3EndpointUrl: ''

## S3 signature version
S3SignatureVersion: 'default'

## S3 Auth version
S3AuthVersion: 'DEFAULT'

## If S3 backend is not Amazon S3 and SSL is enabled on S3 endpoint u
rl then change it to 'True', otherwise keep it as 'False'
S3SslEnabled: False

## If S3 backend is not Amazon S3 and SSL is enabled on S3 endpoint
URL and SSL certificates are self signed, then
## user need to set this parameter value to:
'/etc/nbosdm/s3-cert.pem', otherwise keep it's value
as empty string.
S3SslCert: ''

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

6] Deploy 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 --templates \  
-e /home/stack/templates/node-info.yaml \  
-e /home/stack/templates/overcloud_images.yaml \  
-e /home/stack/nbos-cfg-scripts/redhat-director-scripts/  
  <RHOSP_RELEASE_DIRECTORY>/environments/nbos_env.yaml \  
-e /usr/share/openstack-tripleo-heat-templates/environments/ssl/  
  enable-tls.yaml \  
-e /usr/share/openstack-tripleo-heat-templates/environments/ssl/  
  inject-trust-anchor.yaml \  
-e /home/stack/nbos-cfg-scripts/redhat-director-scripts/<RHOSP_RELEASE_  
  DIRECTORY>/environments/nbos_env_tls_endpoints_public_dns.yaml \  
--ntp-server 192.168.1.34 \  
--libvirt-type qemu \  
--log-file overcloud_deploy.log \  
-r /home/stack/templates/roles_data.yaml
```

7] Verify deployment

Warning: If the containers are in restarting state or not listed by the following command then your deployment is not done correctly. Please recheck if you followed the complete documentation.

7.1] On Controller node

Make sure NetBackup OpenStack `dmapi` and `horizon` containers are in a running state and no other NetBackup OpenStack container is deployed on controller nodes.

When the role for these containers is not **controller**, check on respective nodes according to configured roles_data.yaml.

```
[root@overcloud-controller-0 heat-admin]# podman ps | grep nbos
26fcb9194566  rhosptrainqa.ctlplane.localdomain:8787/nbosdmapl:9.0-rhosp16
kolla_start          5 days ago  Up 5 days ago          nbosdmapl
094971d0f5a9  rhosptrainqa.ctlplane.localdomain:
8787/nbos-horizon-plugin:9.0-rhosp16      kolla_start
5 days ago  Up 5 days ago          horizon
```

7.2] On Compute node

Make sure NetBackup OpenStack datamover container is in the running state and no other NetBackup OpenStack container is deployed on compute nodes.

```
[root@overcloud-novacompute-0 heat-admin]# podman ps | grep nbos
b1840444cc59  rhosptrainqa.ctlplane.localdomain:8787/nbosdm:9.0-rhosp16
kolla_start          5 days ago  Up 5 days ago          nbosdm
```

7.3] On the node with Horizon service

Make sure that horizon container is in the running state. Please note that "Horizon" container is replaced with NetBackup OpenStack Horizon container. This container has the latest OpenStack horizon + NetBackup for OpenStack's horizon plugin.

```
[root@overcloud-controller-0 heat-admin]# podman ps | grep horizon
094971d0f5a9  rhosptrainqa.ctlplane.localdomain:
8787/nbos-horizon-plugin:9.0-rhosp16      kolla_start
5 days ago  Up 5 days ago          horizon
```

8] Additional Steps on NetBackup for OpenStack Appliance

8.1] Change 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)
```

9] Troubleshooting for overcloud deployment failures

NetBackup for OpenStack components are deployed using puppet scripts.

In case of the overcloud deployment failing do the following command to provide the list of errors. The following document also provides valuable insights:

<https://docs.openstack.org/tripleo-docs/latest/install/troubleshooting/troubleshooting-overcloud.html>

```
openstack stack failures list overcloud
heat stack-list --show-nested -f "status=FAILED"
heat resource-list --nested-depth 5 overcloud | grep FAILED
```

=> If nbosdmapi containers does not start well or in restarting state, use following logs to debug.

```
docker logs nbosdmapi
```

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

=> If nbosdm containers does not start well or in restarting state, use following logs to debug.

```
docker logs nbosdm
```

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

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 VM 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.

Details needed for the NetBackup for OpenStack Appliance

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

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 "="
 - 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
 - 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
 - domain the provided user and tenant are located in
 - Format: ID
 - Example: Default
- Administrator
 - User name 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 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 login 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_`

Backup Storage Configuration information

This block requests the necessary information about the backup target that the NetBackup for OpenStack installation will use to store and read backups.

- OpenStack Dist
 - RHOSP requires a special mount point to be used
 - Format: Predefined list
 - Example: RHOSP
- Backup Storage
 - Defines the Backup Storage protocol to use
 - Format: Predefined list of radio buttons
 - Example: NFS

Using the NFS protocol

- NFS Export
 - The path under which the NFS Volumes to be used can be found
 - Format: Comma-separated list of NFS Volumes paths
 - Example: `10.10.2.20:/upstream,10.10.5.100:/nfs2`
- NFS Options
 - NFS options used by the NetBackup for OpenStack Cluster when mounting the NFS Exports
 - Format: NFS options
 - Example: `nolock,soft,timeo=180,intr,lookupcache=none`

Please use the predefined NFS Options and only change them when it is known that changes are necessary.

NetBackup for OpenStack is testing against the predefined NFS options.

Using the S3 protocol

- S3 Compatible
 - Switch between Amazon and other S3 compatible storage solutions
 - Format: Predefined list
 - Example: Amazon S3
- (S3 compatible) Endpoint URL
 - URL to be used to reach and access the provided S3 compatible storage
 - Format: URL
 - Example: objects.nbos.io
- Access Key
 - Access Key necessary to login to the S3 storage
 - Format: access key
 - Example: SFHSAFHPPFSVVBSVBSZRF
- Secret Key
 - Secret Key necessary to login to the S3 storage
 - Format: secret key
 - Example: bfAEURFGHsnvd3435BdfeF
- Region
 - Configured Region for the S3 Bucket (keep the default for S3 compatible without Region)
 - Format: String
 - Example: us-east-1
- Signature Version
 - S3 signature version to use for signing into the S3 storage
 - Format: String
 - Example: Default
- Bucket Name
 - Name of the bucket to be used as Backup target
 - Format: String
 - Example: nbos-backup

Policy Import

Select this box in case of reinitialization or reinstallation of the NetBackup for OpenStack Appliance to import all matching policies that are located on the Backup Target.

Note: Policies that are not assigned to an existing tenant will fail to import and need to be reassigned manually once the configuration is done.

Advanced settings

At the end of the configurator is the option to activate the advanced settings.

Activating this option does provide the possibility to configure the Keystone endpoints that are used for the Datamover API and NetBackup for OpenStack.

Setup NetBackup for OpenStack and Datamover API endpoints.

NetBackup for OpenStack generates Keystone endpoints for two services. The NetBackup for OpenStack Datamover API and the NetBackup for OpenStack Workloadmanager.

Modern OpenStack installation has the endpoint types split over multiple networks. The advanced settings for the Datamover API endpoints and NetBackup for OpenStack Workloadmanager 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 datamover api 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.

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 workloadmgr database, creating the workloadmgr user and setting the right permissions. An example command to create this database would be:

```
create database workloadmgr_auto;  
CREATE USER 'nbos'@'localhost' IDENTIFIED BY 'password';  
GRANT ALL PRIVILEGES ON workloadmgr_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/workloadmgr_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 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, please wait till the configurator finishes. If the configurator does not finish after 6 hours, contact Veritas 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 NBOSVM dashboard to virtual IP.

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:

- `wlm-api`
- `wlm-scheduler`
- `wlm-workloads`

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

```
systemctl status wlm-api
#####
● wlm-api.service - Cluster Controlled wlm-api
   Loaded: loaded (/etc/systemd/system/nbosjm-api.service; disabled;
           vendor preset: disabled)
   Drop-In: /run/systemd/system/wlm-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/wlm-api.service
            └─21265 /home/rhv/myansible/bin/python /usr/bin/workloadmgr-api
              --config-file=/etc/workloadmgr/workloadmgr.conf
```

```
systemctl status wlm-scheduler
#####
● wlm-scheduler.service - Cluster Controlled wlm-scheduler
   Loaded: loaded (/etc/systemd/system/wlm-scheduler.service; disabled;
```

```
vendor preset: disabled)
Drop-In: /run/systemd/system/wlm-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/wlm-scheduler.service
└─21512 /home/rhv/myansible/bin/python /usr/bin/wlm-scheduler
--config-file=/etc/workloadmgr/workloadmgr.conf
```

```
systemctl status wlm-workloads
```

```
#####
```

```
● wlm-workloads.service - workloadmanager workloads service
   Loaded: loaded (/etc/systemd/system/wlm-workloads.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/wlm-workloads.service
           └─20079 /home/rhv/myansible/bin/python
/usr/bin/workloadmgr-workloads
--config-file=/etc/workloadmgr/workloadmgr.conf
           └─20180 /home/rhv/myansible/bin/python
/usr/bin/workloadmgr-workloads
--config-file=/etc/workloadmgr/workloadmgr.conf
           [...]
           └─20181 /home/rhv/myansible/bin/python
/usr/bin/workloadmgr-workloads
--config-file=/etc/workloadmgr/workloadmgr.conf
           └─20233 /home/rhv/myansible/bin/python
/usr/bin/workloadmgr-workloads
--config-file=/etc/workloadmgr/workloadmgr.conf
           └─20236 /home/rhv/myansible/bin/python
/usr/bin/workloadmgr-workloads
--config-file=/etc/workloadmgr/workloadmgr.conf
           └─20237 /home/rhv/myansible/bin/python
/usr/bin/workloadmgr-workloads
--config-file=/etc/workloadmgr/workloadmgr.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.e17_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
wlm-api (systemd:wlm-api): Started om_nbosvm
wlm-scheduler (systemd:wlm-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/
workload_types/detail -X GET -H "X-Auth-Project-Id: admin"
-H "User-Agent: python-workloadmgrclient" -H "Accept:
application/json" -H "X-Auth-Token:
```



```
gAAAAABe40NVFETJeePpk1F9QGGh1LiGnHJVLlgZx9t0HRrK9rC5vq  
KZJRkpAcW1oPH6Q9K9peuHiQrBHEs1-g75Na4xOEESR0LmQJUzP6n3  
7fLfDL_D-hlnjHJZ68iNisIP1fkm9FGSyoyt6IqjO9E7_YVRCTCqNLJ  
67ZkqHuJh1CXwShvjvfw
```

Please check the API guide for more commands and how to generate the X-Auth-Token.

Verify the nbosdm services are up and running

The nbosdm service is the datamover that got installed on all compute nodes. It is recommended to 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.
```

Verify that the NFS Volume is correctly mounted

NetBackup for OpenStack mounts the NFS Backup Target to the NetBackup for OpenStack Appliance and Compute nodes.

To verify those are correctly mounted it is recommended to do the following checks.

First `df -h` looking for `/var/NetBackupOpenStack-mounts/<hash-value>`

```
df -h
#####
Filesystem                Size  Used Avail Use% Mounted on
devtmpfs                  63G   0    63G   0% /dev
tmpfs                     63G   16K   63G   1% /dev/shm
tmpfs                     63G   35M   63G   1% /run
tmpfs                     63G   0    63G   0% /sys/fs/cgroup
/dev/mapper/rhvh-rhvh--
4.3.8.1--0.20200126.0+1
/dev/sda2                 976M  198M  712M  22% /boot
/dev/mapper/rhvh-var      15G   1.9G  12G   14% /var
/dev/mapper/rhvh-home    976M   2.6M  907M   1% /home
/dev/mapper/rhvh-tmp     976M   2.6M  907M   1% /tmp
/dev/mapper/rhvh-var_log  7.8G  230M   7.2G   4% /var/log
/dev/mapper/rhvh-var_log_audit 2.0G   17M   1.8G   1% /var/log/audit
/dev/mapper/rhvh-var_crash 9.8G   37M   9.2G   1% /var/crash
30.30.1.4:/rhv_backup    2.0T   5.3G   1.9T   1% /var/NetBackupOpen
Stack-mounts/MzAuMzAuMS400i9yaHZfYmFja3Vw
30.30.1.4:/rhv_data      2.0T   37G   2.0T   2% /rhev/data-center/
mnt/30.30.1.4:_rhv__data
tmpfs                     13G   0    13G   0% /run/user/0
30.30.1.4:/rhv_iso       2.0T   37G   2.0T   2% /rhev/data-center/
mnt/30.30.1.4:_rhv__iso
```

Secondly do a read, write, and delete test as the user `nova:nova` (`uid = 36 / gid = 36`) from the NetBackup for OpenStack Appliance and the RHV-Host.

```
su nova
#####
[nova@nbosvm MTAuMTAuMi4yMDovdXBzdHJlYW0=]$ touch foo
[nova@nbosvm MTAuMTAuMi4yMDovdXBzdHJlYW0=]$ ll
```

```
total 24
drwxr-xr-x  3 nova nova 4096 Apr  2 17:27 nbosdm_tasks
-rw-r--r--  1 nova nova   0 Apr 23 12:25 foo
drwxr-xr-x  2 nova nova 4096 Apr  2 15:38 test-cloud-id
drwxr-xr-x 10 nova nova 4096 Apr 22 11:00 workload_1540698c-8e22-4dd1-
a898-8f49cd1a898c
drwxr-xr-x  9 nova nova 4096 Apr  8 15:21 workload_51517816-6d5a-4fce-
9ac7-46ee1e09052c
drwxr-xr-x  6 nova nova 4096 Apr 22 11:30 workload_77fb42d2-8d34-4b8d-
bfd5-4263397b636c
drwxr-xr-x  5 nova nova 4096 Apr 23 06:15 workload_85bf16ed-d4fd-49a6-
a753-98c5ca6e906b
[nova@nbosvm MTAuMTAuMi4yMDovdXBzdHJlYW0=]$ rm foo
[nova@nbosvm MTAuMTAuMi4yMDovdXBzdHJlYW0=]$ ll
total 24
drwxr-xr-x  3 nova nova 4096 Apr  2 17:27 nbosdm_tasks
drwxr-xr-x  2 nova nova 4096 Apr  2 15:38 test-cloud-id
drwxr-xr-x 10 nova nova 4096 Apr 22 11:00 workload_1540698c-8e22-4dd1-
a898-8f49cd1a898c
drwxr-xr-x  9 nova nova 4096 Apr  8 15:21 workload_51517816-6d5a-4fce-
9ac7-46ee1e09052c
drwxr-xr-x  6 nova nova 4096 Apr 22 11:30 workload_77fb42d2-8d34-4b8d-
bfd5-4263397b636c
drwxr-xr-x  5 nova nova 4096 Apr 23 06:15 workload_85bf16ed-d4fd-49a6-
a753-98c5ca6e906b
```

Uninstalling from RHOSP

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

- | | |
|---|---|
| Clean NetBackup for OpenStack Datamover API service | See “Clean NetBackup for OpenStack Datamover API service” on page 52. |
| Clean NetBackup for OpenStack Datamover Service | See “Clean NetBackup for OpenStack Datamover Service” on page 53. |
| Clean NetBackup for OpenStack haproxy resources | See “Clean NetBackup for OpenStack haproxy resources” on page 54. |
| Clean NetBackup for OpenStack Keystone resources | See “Clean NetBackup for OpenStack Keystone resources” on page 55. |
| Clean NetBackup for OpenStack database resources | See “Clean NetBackup for OpenStack database resources” on page 55. |

Revert overcloud deploy command	See “Revert overcloud deploy command” on page 56.
Revert back to original RHOSP Horizon container	See “Revert back to original RHOSP Horizon container” on page 56.
Destroy the NetBackup for OpenStack VM Cluster	See “Destroy the NetBackup for OpenStack VM Cluster” on page 56.

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 RHOSP16 onwards
systemctl disable tripleo_nbosdmapi.service
systemctl stop tripleo_nbosdmapi.service
podman stop nbosdmapi
```

Remove `nbosdmapi` container.

```
# For RHOSP16 onwards
podman rm nbosdmapi
podman rm nbosdmapi_init_log
podman rm nbosdmapi_db_sync
```

Clean NetBackup for OpenStack Datamover API service conf directory.

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

Clean NetBackup for OpenStack Datamover API service log directory.

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

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 RHOSP16 onwards
systemctl disable tripleo_nbosdm.service
systemctl stop tripleo_nbosdm.service
podman stop nbosdm
```

Remove `nbosdm` container.

```
# For RHOSP16 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 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 RHOSP16 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 dmapi
openstack user delete dmapi

```

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 RHOSP16
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 dmapi user
MariaDB [mysql]> select user, host from mysql.user where user='dmapi';
+-----+-----+
| user | host          |
+-----+-----+
| dmapi | 172.25.2.10 |

```

```
| dmapi | 172.25.2.8 |
+-----+-----+
2 rows in set (0.00 sec)

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

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

=> Verify that dmapi 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 VM Cluster

List all VMs running on the KVM node

```
virsh list
```

Destroy the NetBackup for OpenStack VMs


```
virsh destroy <NetBackup for OpenStack VM Name or ID>
```

Undefine the NetBackup for OpenStack VMs

```
virsh undefine <NetBackup for OpenStack VM name>
```

Delete the NetBackup for OpenStack VM disk from KVM Host storage

Install workloadmgr CLI client

About the workloadmgr CLI client

The workloadmgr CLI client is provided as rpm and deb packages.

It got tested against the following operating systems:

- CentOS7, CentOS8

Installing the workloadmgr client automatically installs all required OpenStack clients as well.

The installation of the workloadmgr 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.

Installing the workloadmgr client

RPM-based operating systems

The workloadmgr CLI client is available for Python2 and Python3

For Python2 run:

```
yum install nboswlmclient-9.0.999-9.0.noarch.rpm
```

For Python3 run:

```
yum install nboswlmclient-py3-el8-9.0.999-9.0.noarch.rpm
```

Deb-based operating systems

The workloadmgr CLI client is available for Python2 and Python3

For Python2 run:

```
apt-get install nboswlmclient_9.0.999_all.deb
```

For Python3 run:

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

Configuring NetBackup OpenStack Appliance

This chapter includes the following topics:

- [Reconfigure the NetBackup for OpenStack Cluster](#)
- [Configuring the NetBackup master server details](#)
- [Change NetBackup for OpenStack dashboard password](#)
- [Reset NetBackup for OpenStack dashboard password](#)
- [Reinitialize NetBackup for OpenStack](#)
- [Download NetBackup for OpenStack logs](#)

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 master server details

You must configure the master server details on the NetBackup for OpenStack VM. 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 master server details

- 1 Log on to the NetBackup for OpenStack configurator UI.
- 2 Enter the Master server host name.
- 3 Select one of the certificate types.

NBCA

Certificates that the NetBackup CA has issued are referred to as NetBackup CA-signed certificates or NetBackup certificates.

See [“About security management and certificates in NetBackup”](#) on page 61.

External CA

Certificates that are issued by a CA other than the NetBackup CA are referred to as external CA-signed certificates or external certificates.

See [“About security management and certificates in NetBackup”](#) on page 61.

- 4 4. If you select the certificate type NBCA and if the security settings on the master server are configured as **Very High**, you must provide the token.

If the security settings on the master server are the default settings, it works as **Known Hosts** and the token is not required.

5 If you select the certificate type External CA, you must provide the following information:

Certificate file	Provide the path of the certificate file.
Trust store location	Provide the trust store location.
Private key	Provide the private key.
Passphrase file (optional)	Provide the passphrase file if the private key is encrypted.
Use CRL	<p>Select From certificate if you want to use Certificate Revocation List (CRL) defined in the certificate.</p> <p>Select From the following path: if you want to use CRL defined in another file and provide the location of the file.</p> <p>Select Do not use CRL if you do not want to use CRL.</p>

6 Click **Submit**.

7 In the **Ansible Output** tab, you can verify the details such as new certificate on NetBackup OpenStack VM that registers itself as a valid host on the NetBackup Master server.

About security management and certificates in NetBackup

NetBackup uses security certificates to authenticate the NetBackup hosts. These certificates must conform to the X.509 public key infrastructure (PKI) standard. You can use NetBackup certificates or external certificates for secure communication.

NetBackup certificates are issued to hosts by default and the NetBackup master server acts as the CA and manages the Certificate Revocation List (CRL). The NetBackup certificate deployment security level determines how certificates are deployed to NetBackup hosts and how often the CRL is updated on each host. If a host needs a new certificate (the original certificate is expired or revoked), you can use a NetBackup authorization token to reissue the certificate.

External certificates are those that a trusted external CA signed. When you configure NetBackup to use external certificates, the master server, media servers, and clients in the NetBackup domain use the external certificates for secure communication. Additionally, the NetBackup web server uses these certificates for communication between the NetBackup web UI and the NetBackup hosts. Deployment of external certificates, updating or replacing external certificates, and CRL management for the external CA are managed outside of NetBackup.

For more information on external certificates, see the [NetBackup Security and Encryption Guide](#).

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`

Reinitialize NetBackup for OpenStack

The NetBackup for OpenStack Appliance can be reinitialized, which will delete all policy-related values from the NetBackup for OpenStack database.

To reinitialize the NetBackup for OpenStack Appliance do:

- Log on to the NetBackup for OpenStack dashboard
- Click **Admin** in the upper right corner to open the submenu.
- Choose **Reinitialize**.
- Verify that you want to reinitialize the NetBackup for OpenStack.

Download NetBackup for OpenStack logs

It is possible to download the NetBackup for OpenStack logs directly through the NetBackup for OpenStack web GUI.

To download logs through the NetBackup for OpenStack web GUI:

- Log on to the NetBackup for OpenStack web GUI.

- Go to **Logs**.
- Choose the log to be downloaded.
 - Each log for every NetBackup for OpenStack Appliance can be downloaded separately
 - Or a 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:

```
/var/logs/workloadmgr/
```

Configuring NetBackup Master Server

This chapter includes the following topics:

- [License for OpenStack plug-in for NetBackup](#)
- [Allow NetBackup for OpenStack VM on NetBackup master server](#)
- [About launching the OpenStack Horizon UI from the NetBackup web UI](#)

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](#)

Allow NetBackup for OpenStack VM on NetBackup master server

To use the NetBackup for OpenStack VM, you must configure it with the master server. Perform the procedure on the NetBackup master server to allow the communication with NetBackup for OpenStack VM.

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

To allow NetBackup for OpenStack VM on NetBackup master server

- ◆ Run the following command on the NetBackup master server:

- For UNIX

```
bpsetconfig -h masterserver
bpsetconfig> APP_PROXY_SERVER = nbosvm1.domain.org
bpsetconfig> APP_PROXY_SERVER = nbosvm2.domain.org
bpsetconfig> APP_PROXY_SERVER = nbosvm3.domain.org
bpsetconfig>
UNIX systems: <ctl-D>
```

- For Windows

```
bpsetconfig -h masterserver
bpsetconfig> APP_PROXY_SERVER = nbosvm1.domain.org
bpsetconfig> APP_PROXY_SERVER = nbosvm2.domain.org
bpsetconfig> APP_PROXY_SERVER = nbosvm3.domain.org
bpsetconfig>
Windows systems: <ctl-Z>
```

Note: If NetBackup for OpenStack VM is a three-node cluster, you must add all three nodes on the master server.

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

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

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 66.

Table 4-1 Launch OpenStack Horizon UI *(continued)*

Step	Task	Description
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 66.
3	Log on with the role, and launch the Horizon UI.	See “Launching the Horizon UI from the NetBackup web UI” on page 67.

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

NetBackup for OpenStack policies

This chapter includes the following topics:

- [About policies](#)
- [List of policies](#)
- [Create a policy](#)
- [Policy overview](#)
- [Edit a policy](#)
- [Delete a policy](#)
- [Unlock a policy](#)
- [Reset a policy](#)

About policies

A policy is a backup job that protects one or more Virtual Machines according to the configuration. There can be as many policies as needed. But each VM can only be part of one policy.

List of policies

Using Horizon

To view all available policies of a project on Horizon

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

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

- Creation time
- Policy name
- Policy description
- Total snapshots inside this policy
 - Total number of succeeded snapshots
 - Total number of failed snapshots
- Policy type
- Policy status

Using CLI

```
workloadmgr workload-list [--all {True,False}] [--nfsshare <nfsshare>]
```

- `--all {True,False}`List all policies of all projects (valid for admin user only).
- `--nfsshare <nfsshare>`List all policies of NFS share (valid for admin user only).

Create a policy

Using Horizon

To create a policy inside Horizon do the following steps:

- 1 On the Horizon console, navigate to **NBOS Backups > Policies**.
- 2 Click **Create Policy**.
- 3 On the **Details** tab, Provide the policy name, description, and the policy type as Serial or Parallel.
- 4 On the **Policy Members** tab, select the VMs to protect.
- 5 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 backup must repeat.
- 6 Provide the Retention policy on the **Policy Attributes** tab.
- 7 Choose the Full Backup Interval on the **Policy Attributes** tab.

8 If required, select **Pause VM** on the **Options** tab.

9 Click **Create**.

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

Using CLI

```
workloadmgr workload-create --instance <instance-id=instance-uuid>
                             [--display-name <display-name>]
                             [--display-description <display-description>]
                             [--workload-type-id <workload-type-id>]
                             [--source-platform <source-platform>]
                             [--jobschedule <key=key-name>]
                             [--metadata <key=key-name>]
                             [--policy-id <policy_id>]
```

- `--display-name` Optional policy name. (Default=None)
- `--display-description` Optional policy description. (Default=None)
- `--workload-type-id` Policy Type ID is required
- `--source-platform` Policy source platform is required. Supported platform is "OpenStack"
- `--instance` Specify an instance to include in the policy. Specify option multiple times to include multiple instances. Instance-id: Include the instance with this UUID.
- `--jobschedule` Specify following key value pairs for job schedule. Specify option multiple times to include multiple keys. "start_date" : "06/05/2014" "end_date" : "07/15/2014" "start_time" : "2:30 PM" "interval" : "1 hr" "snapshots_to_retain" : "2"
- `--metadata` Specify a key value pair to include in the policy type metadata. Specify option multiple times to include multiple keys. key=value
- `--policy-id <policy_id>` ID of the policy to assign to the policy.

Policy overview

View the information about the policy in the Policy overview.

Using Horizon

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

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy to view.
3. Click the policy name to view the policy overview.

Details	<p>The Policy Details tab provides you with the most important information about the policy:</p> <ul style="list-style-type: none"> ■ Name ■ Description ■ Availability Zone ■ List of protected VMs including the information of qemu guest agent availability <p>The status of the qemu-guest-agent shows whether the necessary OpenStack configuration has been done for this VM to provide qemu guest agent integration. It does not check if the qemu guest agent is installed and configured on the VM.</p> <p>You can navigate to the protected VM directly from the list of protected VMs.</p>
Snapshots	<p>The Snapshots tab shows the list of all available snapshots in the chosen policy.</p> <p>From here it is possible to work with the snapshots, create snapshots on demand and start restores.</p> <p>See "About snapshots" on page 77.</p>
Policy Attributes	<p>The Policy Attributes tab gives an overview of the current configured scheduler and retention policy. The following elements are shown:</p> <ul style="list-style-type: none"> ■ Scheduler Enabled or Disabled ■ Start Date and Time ■ End Date and Time ■ RPO ■ Time until next Snapshot run ■ Retention Policy and Value ■ Full Backup Interval policy and value
File Search	<p>The File Search tab provides access to the powerful search engine, which allows to find files and folders on snapshots without the need of a restore.</p> <p>See "About file search" on page 99.</p>

- Misc. The Miscellaneous tab shows the remaining metadata of the policy. The following information is provided:
- Creation time
 - Last update time
 - Policy ID
 - Policy Type ID
 - Policy Attribute ID
 - Project ID
 - User ID

Using CLI

```
workloadmgr workload-show <workload_id> [--verbose <verbose>]
```

- <workload_id> ID/name of the policy to show.
- --verbose Option to show additional information about the policy.

Edit a policy

Policy can be modified in all components to match changing needs.

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

Using Horizon

To edit a policy in Horizon do the following steps:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy to be modified.
3. Click **Create Snapshot** drop-down.
4. Click **Edit Policy**.
5. Modify the policy as desired. All parameters except policy type can be changed.
6. Click **Update**.

Using CLI

```
usage: workloadmgr  
workload-modify [--display-name <display-name>]
```



```
[--display-description <display-description>]  
[--instance <instance-id=instance-uuid>]  
[--jobschedule <key=key-name>]  
[--metadata <key=key-name>]  
[--policy-id <policy_id>]  
<workload_id>
```

- `--display-name` Optional policy name. (Default=None)
- `--display-description` Optional policy description. (Default=None)
- `--instance <instance-id=instance-uuid>` Specify an instance to include in the policy. Specify option multiple times to include multiple instances. Instance-id: Include the instance with this UUID
- `--jobschedule <key=key-name>` Specify following key value pairs for job schedule Specify option multiple times to include multiple keys. If you don't specify timezone, then by default it takes your local computer timezone
"start_date" : "06/05/2014" "end_date" : "07/15/2014" "start_time" : "2:30 PM"
"interval" : "1 hr" "retention_policy_type" : "Number of Snapshots to Keep" or
"Number of days to retain Snapshots" "retention_policy_value" : "30"
- `--metadata <key=key-name>` Specify a key value pair to include in the policy type metadata Specify option multiple times to include multiple keys. key=value
- `--policy-id <policy_id>` ID of the policy to assign.
- `<workload_id>` ID of the policy to edit.

Delete a policy

Once a policy is no longer needed, it can be safely deleted. All snapshots need to be deleted before the policy gets deleted.

See ["About snapshots"](#) on page 77.

Using Horizon

To delete a policy do the following steps:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy to be deleted.
3. Click **Create Snapshot** drop-down.
4. Click **Delete Policy**.
5. Click **Delete Policy** again to confirm.

Using CLI

```
workloadmgr workload-delete [--database_only <True/False>] <workload_id>
```

- <workload_id> ID/name of the policy to delete.
- --database_only <True/False> Keep True if want to delete from database only. (Default=False)

Unlock a policy

Policies that are actively taking backups or restores are locked for further tasks. You can unlock a policy by force if necessary.

We recommend that you use this feature only as last resort in case of backups or restores being stuck without failing or a restore is required while a backup is running.

Using Horizon

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy to unlock.
3. Click **Create Snapshot** drop-down.
4. Click **Unlock Policy**.
5. Click **Unlock Policy** again to confirm.

Using CLI

```
workloadmgr workload-unlock <workload_id>
```

- <workload_id> ID of the policy to unlock.

Reset a policy

In rare cases it might be necessary to start a backup chain all over again to ensure the quality of the created backups. In case you want to avoid the recreation of the policy, you can reset the policy.

The policy reset will:

- Cancel all ongoing tasks.
- Delete all existing NetBackup for OpenStack snapshots from the protected VMs.
- Recalculate the next snapshot time.

- Take a full backup at the next snapshot.

Using Horizon

To reset a policy do the following steps:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy to reset.
3. Click **Create Snapshot** drop-down.
4. Click **Reset Policy**.
5. Click **Reset Policy** again to confirm.

Using CLI

```
workloadmgr workload-reset <workload_id>
```

- <workload_id> ID/name of the policy to reset.

Performing backups and restores of OpenStack

This chapter includes the following topics:

- [About snapshots](#)
- [List of snapshots](#)
- [Creating a snapshot](#)
- [Snapshot overview](#)
- [Delete snapshots](#)
- [Snapshot Cancel](#)
- [About restores](#)
- [List of Restores](#)
- [Restores overview](#)
- [Delete a Restore](#)
- [Cancel a Restore](#)
- [One-Click Restore](#)
- [Selective Restore](#)
- [Inplace Restore](#)
- [Required restore.json for CLI](#)
- [About file search](#)

- [Navigating to the file search tab in Horizon](#)
- [Configuring and starting a file search in Horizon](#)
- [Start the File Search and retrieve the results in Horizon](#)
- [Doing a CLI File Search](#)
- [About snapshot mount](#)
- [Create a File Recovery Manager Instance](#)
- [Mounting a snapshot](#)
- [Accessing the File Recovery Manager](#)
- [Identifying mounted snapshots](#)
- [Unmounting a snapshot](#)
- [About schedulers](#)
- [Disable a schedule](#)
- [Enable a schedule](#)
- [Modify a schedule](#)
- [About email notifications](#)
- [Requirements to activate email Notifications](#)
- [Activate/Deactivate the email Notifications](#)

About snapshots

A snapshot is a single NetBackup for OpenStack backup of a policy including all data and metadata. It contains the information of all VMs that the policy protects.

List of snapshots

Using Horizon

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy to show the details on.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.

The list of snapshots for the chosen policy contains the following additional information:

- Creation Time
- Name of the snapshot
- Description of the snapshot
- Total amount of Restores from this snapshot
 - Total amount of succeeded Restores
 - Total amount of failed Restores
- Snapshot Type
- Snapshot Size
- Snapshot Status

Using CLI

```
workloadmgr snapshot-list [--workload_id <workload_id>]
                           [--nbos_node <host>]
                           [--date_from <date_from>]
                           [--date_to <date_to>]
                           [--all {True,False}]
```

- `--workload_id <workload_id>` Filter results by workload_id (policy ID).
- `--nbos_node <host>` List all the snapshot operations that are scheduled on a nbos node(Default=None).
- `--date_from <date_from>` From date in format "YYYY-MM-DDTHH:MM:SS" for example 2016-10-10T00:00:00, If you don't specify time then it takes 00:00 by default.
- `--date_to <date_to>` To date in format "YYYY-MM-DDTHH:MM:SS" (default is current day), Specify HH:MM:SS to get snapshots within same day inclusive/exclusive results for date_from and date_to
- `--all {True,False}` List all snapshots 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 case of deactivated scheduler, you can create a snapshot on demand.

Using Horizon

There are two possibilities to create a snapshot on demand.

Possibility 1: From the policy overview

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy to create a snapshot.
3. Click **Create Snapshot**.
4. Provide a name and description for the snapshot.
5. Decide between Full and Incremental Snapshot.
6. Click **Create**.

Possibility 2: From the policy snapshot list

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy to create a snapshot.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.
5. Click **Create Snapshot**.
6. Provide a name and description for the snapshot.
7. Decide between Full and Incremental Snapshot.
8. Click **Create**.

Using CLI

```
workloadmgr workload-snapshot [--full] [--display-name <display-name>]
                                [--display-description <display-description>]
                                <workload_id>
```

- `<workload_id>` ID of the policy to snapshot.
- `--full` Specify if a full snapshot is required.
- `--display-name <display-name>` Optional snapshot name. (Default=None)

- `--display-description <display-description>` Optional snapshot description. (Default=None)

Snapshot overview

Each snapshot contains a lot of information about the backup. This information can be seen in the snapshot overview.

Using Horizon

To reach the snapshot overview follow these steps:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to show.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.
5. Identify the searched snapshot in the snapshot list
6. Click the snapshot name.

Details	<p>The Snapshot Details tab shows the most important information about the snapshot.</p> <ul style="list-style-type: none">■ Snapshot ID/Name/Description■ Snapshot Type■ Time Taken■ Size■ Status■ Which VMs are part of the snapshot■ For each VM in the snapshot<ul style="list-style-type: none">■ Instance Info - Name & Status■ Security Group(s) - Name, Type■ Flavor - vCPUs, Disk, RAM■ Networks - IP, Networkname, Mac Address■ Attached Volumes - Name, Type, size (GB), Mount Point, Restore Size■ Misc - Original ID of the VM
Restores	<p>The Snapshot Restores tab shows the list of restores that have been started from the chosen snapshot. It is possible to start restores from here.</p> <p>See “About restores” on page 83.</p>

- Misc. The Snapshot Miscellaneous tab provides the remaining metadata information about the snapshot.
- Creation Time
 - Last Update time
 - Snapshot ID
 - Policy ID of the Policy containing the snapshot

Using CLI

```
workloadmgr snapshot-show [--output <output>] <snapshot_id>
```

- <snapshot_id> ID of the snapshot to be shown.
- --output <output> Option to get additional snapshot details, Specify --output metadata for snapshot metadata, Specify --output networks for snapshot VMs networks, Specify --output disks for snapshot VMs disks.

Delete snapshots

Once a snapshot is no longer needed, it can be safely deleted from a policy.

The retention policy automatically deletes the oldest snapshots according to the configured policy.

You have to delete all snapshots to be able to delete a policy.

Deleting a NetBackup for OpenStack snapshot does not delete any OpenStack Cinder snapshots. Those need to be deleted separately if desired.

Using Horizon

There are two possibilities to delete a snapshot.

Possibility 1: Single snapshot deletion through the submenu

To delete a single snapshot through the submenu follow these steps:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to delete.
3. Click the policy name to enter the policy overview,
4. Navigate to the **Snapshots** tab.
5. Identify the searched snapshot in the snapshot list
6. Click the drop-down under **Actions** column.

7. Click **Delete Snapshot**.
8. Click **Delete** to confirm.

Possibility 2: Multiple snapshot deletion through checkbox in snapshot overview

To delete one or more snapshots through the snapshot overview do the following:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to show.
3. Click the policy name to enter the policy overview
4. Navigate to the **Snapshots** tab.
5. Identify the searched snapshots in the snapshot list.
6. Select the check box for each snapshot that shall be deleted
7. Click **Delete Snapshots**.
8. Click **Delete** to confirm.

Using CLI

```
workloadmgr snapshot-delete <snapshot_id>
```

- <snapshot_id> ID of the snapshot to be deleted.

Snapshot Cancel

Ongoing snapshots can be canceled.

Canceled snapshots are treated like errored snapshots.

Using Horizon

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to cancel.
3. Click the policy name to enter the policy overview
4. Navigate to the **Snapshots** tab.
5. Identify the searched snapshot in the snapshot list
6. Click **Cancel** on the same line as the identified snapshot.
7. Click **Cancel** to confirm.

Using CLI

```
workloadmgr snapshot-cancel <snapshot_id>
```

- <snapshot_id> ID of the snapshot to be canceled.

About restores

A Restore is the workflow to bring back the backed-up VMs from a NetBackup for OpenStack snapshot.

List of Restores

Using Horizon

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

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to show.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.
5. Identify the searched snapshot in the snapshot list.
6. Click the snapshot name.
7. Navigate to the **Restores** tab.

Using CLI

```
workloadmgr restore-list [--snapshot_id <snapshot_id>]
```

- --snapshot_id <snapshot_id> ID of the snapshot to show the restores of

Restores overview

Using Horizon

To reach the detailed Restore overview follow these steps:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to show.

3. Click the policy name to enter the policy overview
4. Navigate to the **Snapshots** tab.
5. Identify the searched snapshot in the snapshot list.
6. Click the snapshot 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 most important 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 VMs restored
 - Restored VM Name
 - Restored VM Status
 - Restored VM ID

Misc

The Misc tab provides additional Metadata information.

- Creation Time
- Restore ID
- Snapshot ID containing the Restore
- Policy

Using CLI

```
workloadmgr 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 policy.

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

Using Horizon

There are two possibilities to delete a Restore.

Possibility 1: Single Restore deletion through the submenu

To delete a single Restore through the submenu follow these steps:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to delete.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.
5. Identify the searched snapshot in the snapshot list.
6. Click the snapshot 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.

Possibility 2: Multiple Restore deletion through a checkbox in snapshot overview

To delete one or more Restores through the Restore list do the following:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to show.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.
5. Identify the searched snapshot in the snapshot list.
6. Enter the snapshot by clicking the snapshot 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

```
workloadmgr restore-delete <restores_id>
```

- <restore_id> ID of the restore to be deleted

Cancel a Restore

Ongoing Restores can be canceled.

Using Horizon

To cancel a Restore in Horizon follow these steps:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to delete.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.
5. Identify the searched snapshot in the snapshot list.
6. Click the snapshot name.
7. Navigate to the **Restore** tab.
8. Identify the ongoing Restore.
9. Click **Cancel Restore** in the line of the restore in question.
10. Click **Cancel Restore** again to confirm.

Using CLI

```
workloadmgr restore-cancel <restore_id>
```

- <restore_id> ID of the restore to be deleted

One-Click Restore

The One-Click Restore brings back all VMs from the snapshot in the same state as they were backed up. They will:

- be located in the same cluster in the same data center
- use the same storage domain
- connect to the same network
- have the same flavor

The user can't change any Metadata.

The One-Click Restore requires that the original VMs that have been backed up are deleted or otherwise lost. Even if one VM is still running, the One-Click Restore fails.

The One-Click Restore automatically updates the policy to protect the restored VMs.

Using Horizon

There are two possibilities to start a One-click Restore.

Possibility 1: From the snapshot list

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

Possibility 2: From the snapshot overview

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to be restored.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.
5. Identify the snapshot to be restored.
6. Click the snapshot name.

7. Navigate to the **Restores** tab.
8. Click **One-Click Restore**.
9. (Optional) Provide a name/description
10. Click **Create**.

Using CLI

```
workloadmgr snapshot-oneclick-restore [--display-name <display-name>]
                                     [--display-description <display-description>]
                                     <snapshot_id>
```

- `<snapshot_id>` ID of the snapshot to restore.
- `--display-name <display-name>` Optional name for the restore.
- `--display-description <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 VMs to the exact needs of the User.

With the selective restore the following things can be changed:

- Which VMs are getting restored
- Name of the restored VMs
- Which networks to connect with
- Which Storage domain to use
- Which data center or cluster to restore into
- Which flavor the restored VMs will use

The Selective Restore is always available and doesn't have any prerequisites.

Using Horizon

There are two possibilities to start a Selective Restore.

Possibility 1: From the snapshot list

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to be restored.
3. Click the policy name to enter the policy overview.

4. Navigate to the **Snapshots** tab.
5. Identify the snapshot to be restored.
6. From the drop-down menu under **Actions** column, select **Selective Restore**.
7. Configure the Selective Restore as desired.
8. Click **Restore**.

Possibility 2: From the snapshot overview

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to be restored.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.
5. Identify the snapshot to be restored.
6. Click the snapshot name.
7. Navigate to the **Restores** tab.
8. Click **Selective Restore**.
9. Configure the selective Restore as desired.
10. Click **Restore**.

Using CLI

```
workloadmgr snapshot-selective-restore [--display-name <display-name>]
                                     [--display-description <display-description>]
                                     [--filename <filename>]
                                     <snapshot_id>
```

- `<snapshot_id>` ID of the snapshot to restore.
- `--display-name <display-name>` Optional name for the restore.
- `--display-description <display-description>` Optional description for restore.
- `--filename <filename>` Provide file path(relative or absolute) including file name , by default it will read file:
`/usr/lib/python2.7/site-packages/workloadmgrclient/input-files/restore.json` . You can use this for reference or replace values into this file.

Inplace Restore

The Inplace Restore covers those use cases, where the VM 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 Inplace Restore only works when the original VM and the original volume are still available and connected. NetBackup for OpenStack is checking this by the saved Object-ID.

The Inplace Restore will not create any new RHV resources. Please use one of the other restore options if new volumes or VMs are required.

Using Horizon

There are two possibilities to start an Inplace Restore.

Possibility 1: From the snapshot list

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to be restored.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.
5. Identify the snapshot to be restored.
6. From the drop-down under **Actions** column, select **Inplace Restore**.
7. Configure the Inplace Restore as desired.
8. Click **Restore**.

Possibility 2: From the snapshot overview

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to be restored.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.
5. Identify the snapshot to be restored.
6. Click the snapshot name.
7. Navigate to the **Restores** tab.
8. Click **Inplace Restore**.
9. Configure the Inplace Restore as desired.

10. Click **Restore**.

Using CLI

```
workloadmgr snapshot-inplace-restore [--display-name <display-name>]
                                     [--display-description <display-description>]
                                     [--filename <filename>]
                                     <snapshot_id>
```

- `<snapshot_id>` ID of the snapshot to restore.
- `--display-name <display-name>` Optional name for the restore.
- `--display-description <display-description>` Optional description for restore.
- `--filename <filename>` Provide file path(relative or absolute) including file name , by default it will read file:
`/usr/lib/python2.7/site-packages/workloadmgrclient/input-files/restore.json` .You can use this for reference or replace values into this file.

Required restore.json for CLI

The workloadmgr client CLI is using a restore.json file to define the restore parameters for the selective and the inplace 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 snapshot overview.

```
{
  name: getjson,
  description: -,
  oneclickrestore: False,
  restore_type: selective,
  type: openstack,
  openstack:
    {
      instances:
        [
          {
            include: True,
```

```
id: 890888bc-a001-4b62-a25b-484b34ac6e7e,  
name: cdcentOS-1,  
availability_zone:,  
nics: [],  
vdisks:  
  [  
    {  
      id: 4cc2b474-1f1b-4054-a922-497ef5564624,  
      new_volume_type:,  
      availability_zone: nova  
    }  
  ],  
flavor:  
  {  
    ram: 512,  
    ephemeral: 0,  
    vcpus: 1,  
    swap:,  
    disk: 1,  
    id: 1  
  }  
},  
restore_topology: True,  
networks_mapping:  
  {  
    networks: []  
  }  
}
```

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 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.

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 snapshot to restore and how they are to be restored.

Even when VMs 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 boot 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'
}
```

Inplace Restore required information

The Inplace 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 VM 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.
- `vdisks` 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 empty value for the restore to work.

Full Inplace 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',
```

```
        'vdisks':[
            {
                'restore_cinder_volume':True,
                'id':'04d66b70-6d7c-4d1b-98e0-11059b89cba6',
                'new_volume_type':'ceph'
            }
        ]
    },
    ],
    'restore_topology':False,
    'networks_mapping':{
        'networks':[
        ]
    }
}
}
```

About file search

The file search functionality let's you search for files and folders that are located on a selected VM in a policy in one or more backups.

Navigating to the file search tab in Horizon

The file search tab is part of every policy overview. To reach it follow these steps:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy a file search shall be done in.
3. Click the policy name to enter the policy overview.
4. Click **File Search** to enter the file search tab.

Configuring and starting a file search in Horizon

A file search runs against a single virtual machine for a chosen subset of backups using a provided search string.

To run a file search the following elements need to be decided and configured.

Choose the VM the file search shall run against

Under VM Name/ID, choose the VM that the search is done upon. The drop-down menu provides a list of all VMs that are part of any snapshot in the policy.

VMs that are no longer actively protected by the policy but are still part of an existing snapshot are listed in red.

Set the file path

The file path defines the search string that is run against the chosen VM and snapshots. This search string does support basic RegEx.

The file path has to start with a "/".

Windows partitions are fully supported. Each partition is its own volume with its own root. Use "/Windows" instead of "C:"

The file search does not go into deeper directories and always searches on the directory provided in the file path.

Example file path for all files inside /etc : /etc/*

Define the Snapshots to search in

"Filter Snapshots by" is the third and last component that needs to be set. This defines which snapshots are going to be searched.

There are 3 possibilities for a pre-filtering:

1. All Snapshots: Lists all Snapshots that contain the chosen VM from all available snapshots
2. Last Snapshots: Choose between the last 10, 25, 50, or custom Snapshots and click Apply to get the list of the available Snapshots for the chosen VM that match the criteria.
3. Date Range: Set a start and end date and click apply to get the list of all available Snapshots for the chosen VM within the set dates.

After the pre-filtering is done, all matching snapshots are automatically pre-selected. Uncheck any snapshot that shall not be searched.

Note: When no snapshot is chosen the file search will not start.

```
[--start_filter <start_filter>]  
[--date_from <date_from>]  
[--date_to <date_to>]  
<vm_id> <file_path>
```

- <vm_id> ID of the VM to be searched
- <file_path> Path of the file to search for
- --snapshotids <snapshotid> Search only in specified snapshot ids
snapshot-id: include the instance with this UUID
- --end_filter <end_filter> Displays last snapshots, example , last 10
snapshots, default 0 means displays all snapshots
- --start_filter <start_filter> Displays snapshots starting from , example
, snapshot starting from 5, default 0 means starts from first snapshot
- --date_from <date_from> From date in format "YYYY-MM-DDTHH:MM:SS"
eg 2016-10-10T00:00:00, If time isn't specified then it takes 00:00 by default
- --date_to <date_to> To date in format "YYYY-MM-DDTHH:MM:SS"(default is
current day),Specify HH:MM:SS to get snapshots within same day
inclusive/exclusive results for date_from and date_to

About snapshot mount

NetBackup for OpenStack allows you to view or download a file from the snapshot. Any changes to the files or directories when snapshot is mounted are temporary and are discarded when the snapshot is unmounted. Mounting is a faster way to restore a single or multiple files. To mount a snapshot follow these steps.

Create a File Recovery Manager Instance

It is recommended to do these steps once to the chosen cloud-Image and then upload the modified cloud image to Glance.

- Create an OpenStack image using a Linux based cloud-image like Ubuntu, CentOS or RHEL with the following metadata parameters.

```
openstack image create \  
--file <File Manager Image Path> \  
--container-format bare \  
--disk-format qcow2 \  
--public \  

```

```
--property hw_qemu_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. Bigger snapshots can require more RAM.

Steps to apply on CentOS and RHEL cloud-images

- Install and activate qemu-guest-agent.
- Edit `/etc/sysconfig/qemu-ga` and remove the following from `BLACKLIST_RPC` section

```
guest-file-read  
guest-file-write  
guest-file-open  
guest-file-close
```

- Disable SELINUX in `/etc/sysconfig/selinux`.

```
SELINUX=disabled
```

- Install python3.

```
yum install python3
```

- Install lvm2

```
yum install lvm2
```

- Restart the Instance.

Mounting a snapshot

Mounting a snapshot to a File Recovery Manager provides read access to all the data that is located in the mounted snapshot.

Unmount any mounted snapshot once there is no further need to keep it mounted. Mounted snapshots will not be purged by the Retention policy.

It is possible to run the mounting process against any OpenStack instance. During this process the instance is restarted.

Always mount snapshots to File Recovery Manager instances only.

Using Horizon

There are 2 possibilities to mount a snapshot in Horizon.

Through the snapshot list

To mount a snapshot through the snapshot list follow these steps:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to mount.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.
5. Identify the searched snapshot in the snapshot list.
6. From the drop-down under **Actions** column, select **OneClick Restore**.
7. Click **Mount Snapshot**.
8. Choose the File Recovery Manager instance to mount to.
9. Click **Mount** to confirm.

Should all instances of the project be listed and there is a File Recovery Manager instance existing verify together with the administrator that the File Recovery Manager image has the following property set:

```
nbos_recovery_manager=yes
```

Through the File Search results

To mount a snapshot through the File Search results follow these steps:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to mount.
3. Click the policy name to enter the policy overview.
4. Navigate to the **File Search** tab.
5. Do a File Search.
6. Identify the snapshot to be mounted.
7. Click **Mount Snapshot** for the chosen snapshot.
8. Choose the File Recovery Manager instance to mount to.
9. Click **Mount**.

Should all instances of the project be listed and there is a File Recovery Manager instance existing verify together with the administrator that the File Recovery Manager image has the following property set:

```
nbos_recovery_manager=yes
```

Using CLI

```
workloadmgr snapshot-mount <snapshot_id> <mount_vm_id>
```

- <snapshot_id> ID of the snapshot to be mounted
- <mount_vm_id> ID of the File Recovery Manager instance to mount the snapshot to.

Accessing the File Recovery Manager

The File Recovery Manager is a normal Linux based OpenStack instance.

It can be accessed via SSH or SSH based tools like FileZilla or WinSCP.

SSH Log in is often disabled by default in cloud-images. Enable SSH Log in if necessary.

The mounted snapshot can be found at the following path:

```
/home/ubuntu/nbos-mounts/mounts/
```

Each VM in the snapshot has its own directory using the VM_ID as the identifier.

Identifying mounted snapshots

Sometimes a snapshot is mounted for a longer duration and hence it is important to be identified.

Using Horizon

There are 2 possibilities to identify mounted snapshots 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 snapshot ID of the snapshot that has been mounted last.

Note: This value only gets updated, when a new snapshot is mounted.

From the snapshot list

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to mount.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.
5. Search for the snapshot that has the option **Unmount Snapshot**.

Using CLI

```
workloadmgr snapshot-mounted-list [--workloadid <workloadid>]
```

- `--workloadid <workloadid>` Restrict the list to snapshots in the provided policy

Unmounting a snapshot

Once a mounted snapshot is no longer needed it is possible and recommended to unmount the snapshot.

Unmounting a snapshot frees the File Recovery Manager instance to mount the next snapshot and allows NetBackup for OpenStack retention policy to purge the former mounted snapshot.

Warning: Deleting the File Recovery Manager instance does not update the NetBackup for OpenStack appliance. The snapshot will be considered mounted until an unmount command has been received.

Using Horizon

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy that contains the snapshot to mount.
3. Click the policy name to enter the policy overview.
4. Navigate to the **Snapshots** tab.
5. Search for the snapshot that has the option **Unmount Snapshot**.

6. Click **Unmount Snapshot**.

Using CLI

```
workloadmgr snapshot-dismount <snapshot_id>
```

- <snapshot_id> ID of the snapshot to unmount.

About schedulers

Every policy 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

Disable a schedule

Using Horizon

To disable the scheduler of a single policy in Horizon do the following steps:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy to be modified.
3. From the drop-down under **Actions** column, select **Edit Policy**.
4. Navigate to the tab **Schedule**.
5. Clear **Enabled**.
6. Click **Update**.

Using CLI

```
workloadmgr disable-scheduler --workloadids <workloadid>
```

- --workloadid <workloadid> Requires at least one policy ID, Specify an ID of the policy whose scheduler disables. Specify option multiple times to include multiple policies. --workloadids <workloadid> --workloadids <workloadid>

Enable a schedule

Using Horizon

To enable the scheduler of a single policy in Horizon do the following steps:

1. On the Horizon console, navigate to **NBOS Backups > Policies**.
2. Identify the policy to be modified.
3. From the drop-down under **Actions** column, select **Edit Policy**.
4. Navigate to the tab **Schedule**.
5. Select **Enabled**.
6. Click **Update**.

Using CLI

```
workloadmgr enable-scheduler --workloadids <workloadid>
```

- `--workloadid <workloadid>` Requires at least one policy id, Specify an ID of the policy whose scheduler enables. Specify option multiple times to include multiple policies. `--workloadids <workloadid> --workloadids <workloadid>`

Modify a schedule

To modify a schedule the policy itself needs to be modified.

See [“Edit a policy”](#) on page 72.

About email notifications

NetBackup for OpenStack does provide the possibility to notify users via email after every backup and restore.

The email will be send to the owner of the policy.

Requirements to activate email Notifications

To use the email notifications 2 requirements need to be met.

Both requirements need to be set or configured by the OpenStack Administrator. Please contact your OpenStack Administrator to verify the requirements.

- User email assigned

As the email will be sent to the owner of the policy the OpenStack User, who created the policy, is required to have an email address associated.

- NetBackup for OpenStack E-Mail Server configured
NetBackup for OpenStack needs to know which E-Mail server to use to send the email notifications. Backup Administrators can do this in their specific area inside Horizon.

Activate/Deactivate the email Notifications

Email notifications are activated tenant wide. To activate the email notification feature for a tenant follow these steps:

1. On the Horizon console, navigate to **NBOS Backups > Settings**.
2. Check/Uncheck the box for **Enable Email Alerts**.

Performing Backup Administration tasks

This chapter includes the following topics:

- [NBOS Backup Admin Area](#)
- [Policy Attributes](#)
- [Policy Quotas](#)
- [Managing Trusts](#)
- [Policy import and migration](#)
- [Disaster Recovery](#)
- [Example runbook for disaster recovery using NFS](#)

NBOS Backup Admin Area

NetBackup for OpenStack provides Backup as a Service, which allows OpenStack users to manage and control their backups themselves. This doesn't eradicate the need for a Backup Administrator, who has an overview of the complete backup solution.

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 CLI.

Access the NBOS Backup Admin area

To access the NBOS Backup Admin area follow these steps:

1. Log on to Horizon using admin user.

2. Navigate to **Admin > NBOS Backup Admin > NetBackupOpenStack**.

The **NBOS Backup Admin** area provides the following features.

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 most needed information at a glance, including:

- Storage Usage (NFS only)
- Number of protected VMs compared to number of existing VMs
- 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.

Policies tab

This tab provides information about all currently existing policies. 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 policy
- Project that contains the policy
- Policy name
- Policy type
- Availability Zone
- Number of protected VMs
- 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 (Blue) Backups compared to Incremental (Red) Backups
- Number of successful backups
- Number of failed backups

- Storage that is used by that policy
- Which backup target is used
- When is the next snapshot run
- What is the general interval of the policy
- Scheduler status including a switch to deactivate/activate the policy

Usage tab

Administrators often need to figure out, where a lot of resources are used up, 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
- VMs protected by a Tenant

It is possible to drill down to see the same information per policy and finally per protected VM.

The Usage tab includes policies and VMs 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 it's own node. It is typically shown directly below the current active Controller Node.

NBOSDM tab (NetBackup for OpenStack Datamover Service)

This tab displays information about NetBackup for OpenStack Datamover service. The following information is shown:

- Service-Name

- Compute node the service is running on
- Zone
- Service Status from OpenStack perspective (enabled or disabled)
- Version of the Service
- General Status
- Last time the Status was updated

Storage tab

This tab displays information about the backup target storage. It contains the following information:

- Storage Name

Clicking on the Storage name provides an overview of all policies that are stored on that storage.

- Capacity of the storage
- Total utilization of the storage
- Status of the storage
- Statistic information
 - Percentage all storages are used
 - Percentage how much storage is used for full backups
 - Number of Full backups versus Incremental backups

Audit tab

Audit logs provide the sequence of policy-related activities that users perform such as policy creation, snapshot creation, and so on. The following information is shown:

- Date and time of the entry
- What task has been done
- Project the task has performed in
- User that performed the task

The Audit log can be searched for strings to find for example only entries done by a specific user.

Additionally, the shown time frame can be changed as necessary.

Policy Attributes tab

The Policy Attributes tab gives administrators the possibility to work with policy attributes.

See [“Policies tab”](#) on page 111.

Settings tab

This tab manages all global settings for the whole cloud. NetBackup for OpenStack has two types of settings:

1. Email settings
2. Job scheduler settings.

Email Settings

These settings are used by NetBackup for OpenStack to send email reports of snapshots and restores to users.

Configuring the Email settings is a must-have 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 CLI use the following commands:

To set an email setting for the first time or after deletion use:

```
workloadmgr 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 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.

To update an already set email setting through CLI use:

```
workloadmgr 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.

To show an already set email setting use:

```
workloadmgr 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.

To delete a set email setting use:

```
workloadmgr setting-delete <setting_name>
```

- `<setting_name>` Name of the setting to delete.

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

Disable or enable Job Scheduler

The Global Job Scheduler can be used to deactivate all scheduled policies without modifying each one of them.

To disable or enable the Global Job Scheduler through the Backups-Admin area:

1. Log on to Horizon using admin 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.

The Global Job Scheduler can be controlled through CLI as well.

To get the status of the Global Job Scheduler, use:

```
workloadmgr get-global-job-scheduler
```

To deactivate the Global Job Scheduler, use:

```
workloadmgr disable-global-job-scheduler
```

To activate the Global Job Scheduler, use:

```
workloadmgr enable-global-job-scheduler
```

Policy Attributes

NetBackup for OpenStack's tenant driven backup service gives tenants control over backup policies. However, sometimes it may be too much control to tenants and the cloud administrators may want to limit what policies 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 backup policy, it may affect the resource limits set on the cloud infrastructure. Instead, if the cloud administrator can define predefined backup policies and each tenant is only limited to those policies then cloud administrators can exert better control over backup service.

Policy is similar to nova flavor where a tenant cannot create arbitrary instances. Instead, each tenant is only allowed to use the nova flavors published by the admin.

List the available policies

Using Horizon

To see all available policies in Horizon follow these steps:

1. Log on to Horizon using admin user.
2. Navigate to **Admin > NBOS Backup Admin > NetBackupOpenStack > Policy Attributes**.

The following information is shown in the policy tab for each available policy:

- Creation time
- Name
- Description
- Status
- Set interval

- Set retention type
- Set retention value
- Full Backup Interval
- Action

Using CLI

```
workloadmgr policy-list
```

```
workloadmgr policy-show <policy_id>
```

- <policy_id> ID of the policy to show.

Create policy attributes

Using Horizon

To create policy attributes in Horizon follow these steps:

1. Log on to Horizon using admin user.
2. Navigate to **Admin > NBOS Backup Admin > NetBackupOpenStack > Policy Attributes**.
3. Click **New Policy Attributes**.
4. Provide the policy attribute name and the description on the **Details** tab.
5. Provide the RPO on the **Policy** tab.
6. Select the Snapshot Retention Type and the value.
 - Number of snapshots to keep
 - Number of days to retain snapshots
7. Select one of the options for Full Backup Interval.
 - Never
 - Always
 - Number of Snapshots
8. Click **Create**.

Using CLI

```
workloadmgr policy-create --policy-fields <key=key-name>
                        [--display-description <display_description>]
                        [--metadata <key=key-name>]
                        <display_name>
```

- `--policy-fields <key=key-name>` Specify following key value pairs for policy fields. Specify option multiple times to include multiple keys. "interval" : "1 hr" "retention_policy_type" : "Number of Snapshots to Keep" or "Number of days to retain Snapshots" "retention_policy_value" : "30" "fullbackup_interval" : "-1" (Enter Number of incremental snapshots to take Full Backup between 1 to 999, "-1" for "NEVER" and "0" for "ALWAYS")For example `--policy-fields interval="1 hr" --policy-fields retention_policy_type="Number of Snapshots to Keep"--policy-fields retention_policy_value="30" --policy- fields fullbackup_interval="2"`
- `--display-description <display_description>` Optional policy description. (Default=No description)
- `--metadata <key=keyname>` Specify key value pairs to include in the policy type metadata. Specify option multiple times to include multiple keys. `key=value`
- `<display_name>` The name the policy will get.

Edit the policy attribute

Using Horizon

To edit a policy in Horizon follow these steps:

1. Log on to Horizon using admin user.
2. Navigate to **Admin > NBOS Backup Admin > NetBackupOpenStack > Policy Attributes**.
3. Identify the policy to edit.
4. Click **Edit Policy Attribute** at the end of the line of the chosen policy.
5. Edit the policy attribute as desired - all values can be changed.
6. Click **Update**.

Using CLI

```
workloadmgr policy-update [--display-name <display-name>]
                        [--display-description <display-description>]
```

```

[--policy-fields <key=key-name>]
[--metadata <key=key-name>]
<policy_id>
    
```

- --display-name <display-name> Name of the policy attribute.
- --display-description <display_description> Optional policy attribute description. (Default=No description)
- --policy-fields <key=key-name> Specify following key value pairs for policy fields. Specify option multiple times to include multiple keys. "interval" : "1 hr" "retention_policy_type" : "Number of Snapshots to Keep" or "Number of days to retain Snapshots" "retention_policy_value" : "30" "fullbackup_interval" : "-1" (Enter Number of incremental snapshots to take Full Backup between 1 to 999, "-1" for "NEVER" and "0" for "ALWAYS") For example --policy-fields interval="1 hr" --policy-fields retention_policy_type="Number of Snapshots to Keep"--policy-fields retention_policy_value="30" --policy- fields fullbackup_interval="2"
- --metadata <key=keyname> Specify a key value pair to include in the policy type metadata. Specify option multiple times to include multiple keys. key=value
- <policy_id> The name the policy will get.

Assign/Remove a policy

Using Horizon

To assign or remove a policy in Horizon follow these steps:

1. Log on to Horizon using admin user.
2. Navigate to **Admin > NBOS Backup Admin > NetBackupOpenStack > Policy Attributes**.
3. Identify the policy attribute to assign/remove.
4. Click the drop-down at the end of the line of the chosen policy.
5. Click **Add/Remove Projects**.
6. Choose projects to add or remove by using the plus or minus options.
7. Click **Apply**.

Using CLI

```
workloadmgr policy-assign [--add_project <project_id>]
```



```
[--remove_project <project_id>]  
<policy_id>
```

- `--add_project <project_id>` ID of the project to assign policy to. Use multiple times to assign multiple projects.
- `--remove_project <project_id>` ID of the project to remove policy from. Use multiple times to remove multiple projects.
- `<policy_id>` Policy to be assigned or removed

Delete a policy

Using Horizon

To delete a policy in Horizon follow these steps:

1. Log on to Horizon using admin user.
2. Navigate to **Admin > NBOS Backup Admin > NetBackupOpenStack > Policy Attributes**.
3. Identify the policy to assign/remove.
4. Click the drop-down at the end of the line of the chosen policy.
5. Click **Delete Policy**
6. Click **Delete** to confirm.

Using CLI

```
workloadmgr policy-delete <policy_id>
```

- `<policy_id>` ID of the policy to be deleted.

Policy Quotas

NetBackup for OpenStack enables OpenStack administrators to set Project Quotas against the usage of NetBackup for OpenStack.

The following Quotas can be set:

- Number of policies a Project is allowed to have
- Number of Snapshots a Project is allowed to have
- Number of VMs a Project is allowed to protect
- Amount of Storage a Project is allowed to use on the Backup Target


Work with Policy Quotas via Horizon

The NetBackup for OpenStack Quota feature is available for all supported OpenStack versions and distributions, but only Train and higher releases include the Horizon integration of the Quota feature.

Policy Quotas are managed like any other Project Quotas.

1. Log on to Horizon using admin user.
2. Navigate to **Identity > Projects**.
3. Identify the project to modify or show the quotas on.
4. Click the drop-down at the end of the line of the chosen project.
5. Click **Modify Quotas**.
6. Navigate to **Policy Manager**.
7. Edit Quotas as desired.
8. Click **Save**.

Figure 7-1 Horizon integration for Policy Manager Quotas



The screenshot displays the 'Edit Quotas' window in the Horizon interface. At the top right, there is a close button (X). Below the title, there are four tabs: 'Compute', 'Volume', 'Network', and 'NBOSJM'. The 'Compute' tab is active. Underneath the tabs, there are four rows of input fields, each with a label and a value of '-1':

Category	Value
Policies	-1
Snapshots	-1
Protected VMs	-1
Storage (Bytes)	-1

At the bottom right of the form, there are two buttons: 'Cancel' and 'Save'.

Work with Policy Quotas via CLI

List available Quota Types

NetBackup for OpenStack is providing several different Quotas. The following command allows listing those.

Note: NetBackup for OpenStack 9.0.0.1 does not yet have the Quota Type Volume integrated. Using this will not generate any Quotas a Tenant has to apply to.

```
workloadmgr project-quota-type-list
```

Show Quota Type Details

The following command will show the details of a provided Quota Type.

```
workloadmgr project-quota-type-show <quota_type_id>
```

- <quota_type_id> ID of the Quota Type to show.

Create a Quota

The following command will create a Quota for a given project and set the provided value.

```
workloadmgr project-allowed-quota-create --quota-type-id quota_type_id
                                         --allowed-value allowed_value
                                         --high-watermark high_watermark
                                         --project-id project_id
```

- <quota_type_id> ID of the Quota Type to be created.
- <allowed_value> Value to set for this Quota Type.
- <high_watermark> Value to set for High Watermark warnings.
- <project_id> Project to assign the quota to.

The high watermark is automatically set to 80% of the allowed value when set via Horizon.

A created Quota will generate an `allowed_quota_object` with its own ID. This ID is needed when continuing to work with the created Quota.

List allowed Quotas

The following command lists all NetBackup for OpenStack Quotas set for a given project.

```
workloadmgr project-allowed-quota-list <project_id>
```

- `<project_id>` Project to list the Quotas from.

Show allowed Quota

The following command shows the details about a provided allowed Quota.

```
workloadmgr project-allowed-quota-show <allowed_quota_id>
```

- `<allowed_quota_id>` ID of the allowed Quota to show.

Update allowed Quota

The following command shows how to update the value of an already existing allowed Quota.

```
workloadmgr project-allowed-quota-update [--allowed-value <allowed_value>]
                                         [--high-watermark <high_watermark>]
                                         [--project-id <project_id>]
                                         <allowed_quota_id>
```

- `<allowed_value>` Value to set for this Quota Type.
- `<high_watermark>` Value to set for High Watermark warnings.
- `<project_id>` Project to assign the quota to.
- `<allowed_quota_id>` ID of the allowed Quota to update.

Delete allowed Quota

The following command will delete an allowed Quota and sets the value of the connected Quota Type back to unlimited for the affected project.

```
workloadmgr project-allowed-quota-delete <allowed_quota_id>
```

- `<allowed_quota_id>` ID of the allowed Quota to delete.

Managing Trusts

NetBackup for OpenStack is using 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 policy.

Trusts can only be worked with via CLI

List all trusts

```
workloadmgr trust-list
```

Show a trust

```
workloadmgr trust-show <trust_id>
```

- <trust_id> ID of the trust to show.

Create a trust

```
workloadmgr trust-create [--is_cloud_trust {True,False}] <role_name>
```

- <role_name> Name of the role that trust is created for.
- --is_cloud_trust {True,False} Set to true if creating 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.

Delete a trust

```
workloadmgr 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 secures, 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.

Import policies

Policy import allows to import policies existing on the Backup Target into the NetBackup for OpenStack database.

The policy import is designed to import policies, which are owned by the Cloud. It will not import or list any policies that are owned by a different cloud.

To get a list of importable policies use the following CLI command:

```
workloadmgr workload-get-importworkloads-list [--project_id <project_id>]
```

- `--project_id <project_id>` List policies that belong to given project only.

To import policies into the NetBackup for OpenStack database use the following CLI command:

```
workloadmgr workload-importworkloads [--workloadids <workloadid>]
```

- `--workloadids <workloadid>` Specify policy ids to import. Repeat option for multiple policies.

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:

```
workloadmgr workload-get-orphaned-workloads-list [--migrate_cloud  
{True,False}]  
[--generate_yaml {True,False}]
```

- `--migrate_cloud {True,False}` Set to True if you want to list policies from other clouds as well. Default is False.
- `--generate_yaml {True,False}` Set to True if you want to generate output file in yaml format, which would be further used as input for policy reassign API.

Running this command against a Backup Target with many policies can take a bit of time. NetBackup for OpenStack is reading the complete Storage and verifies every found policy against the policies known in the database.

Reassigning policies

OpenStack administrators are able to reassign a policy to a new owner. This involves the possibility to migrate a policy from one cloud to another or between projects.

Warning: Reassigning a policy only changes the database of the target NetBackup for OpenStack installation. Now that it is managed by a different NetBackup installation, the original source installation is not updated.

Use the following CLI command to reassign a policy:

```
workloadmgr workload-reassign-workloads  
[--old_tenant_ids <old_tenant_id>]  
[--new_tenant_id <new_tenant_id>]  
[--workload_ids <workload_id>]  
[--user_id <user_id>]  
[--migrate_cloud {True,False}]  
[--map_file <map_file>]
```

- `--old_tenant_ids <old_tenant_id>` Specify old tenant ids from which policies need to reassign to new tenant. Specify multiple times to choose policies from multiple tenants.
- `--new_tenant_id <new_tenant_id>` Specify new tenant id to which policies need to reassign from old tenant. Only one target tenant can be specified.

- `--workload_ids <workload_id>` Specify policy ids which need to reassign to new tenant. If not provided then all the policies from old tenant will get reassigned to new tenant. Specify multiple times for multiple policies.
- `--user_id <user_id>` Specify user id to which policies need to reassign from old tenant. only one target user can be specified.
- `--migrate_cloud {True,False}` Set to True if you want to reassign policies from other clouds as well. Default if False
- `--map_file` Provide file path(relative or absolute) including file name of reassign map file. Provide list of old policies mapped to new tenants. Format for this file is YAML.

A sample mapping file with explanations is shown below:

```
reassign_mappings:
- old_tenant_ids: [] #user can provide list of old_tenant_ids or
  workload_ids
  new_tenant_id: new_tenant_id
  user_id: user_id
  workload_ids: [] #user can provide list of old_tenant_ids or workload_ids
  migrate_cloud: True/False #Set to True if want to reassign workloads from
  # other clouds as well. Default is False

- old_tenant_ids: [] #user can provide list of old_tenant_ids or
  workload_ids
  new_tenant_id: new_tenant_id
  user_id: user_id
  workload_ids: [] #user can provide list of old_tenant_ids or workload_ids
  migrate_cloud: True/False #Set to True if want to reassign workloads from
  # other clouds as well. Default is False
```

Disaster Recovery

NetBackup for OpenStack policies are designed to allow a Disaster Recovery without the need to backup the NetBackup for OpenStack database.

As long as the NetBackup for OpenStack policies are existing on the Backup Target Storage and a NetBackup for OpenStack installation has access to them, it is possible to restore the policies.

Disaster Recovery Process

1. Install and Configure NetBackup for OpenStack for the target cloud.
See [“Installing NetBackup for OpenStack Components”](#) on page 28.
2. Verify required mount-paths and create if necessary.
See [“Mount-paths”](#) on page 129.
3. Reassign policies.
See [“Reassigning policies”](#) on page 127.
4. Notify users of policies being available.

This procedure is designed to be applicable to all OpenStack installations using NetBackup for OpenStack. It is to be used as a starting point to develop the exact Disaster Recovery process of a specific environment.

Instead of notifying the users, the policy should be restored as necessary to have the user in each project that has the necessary privileges to restore.

Mount-paths

NetBackup for OpenStack incremental Snapshots involve a backing file to the prior backup taken, which makes every NetBackup for OpenStack incremental backup a synthetic full backup.

NetBackup for OpenStack is using qcow2 backing files for this feature:

```
qemu-img info 85b645c5-c1ea-4628-b5d8-1faea0e9d549
image: 85b645c5-c1ea-4628-b5d8-1faea0e9d549
file format: qcow2
virtual size: 1.0G (1073741824 bytes)
disk size: 21M
cluster_size: 65536
backing file: /var/NetBackupOpenStack-mounts/MTAuMTAuMi4yMDovdXBzdHJlYW0=
/workload_3c2fbee5-ad90-4448-b009-5047bcffc2ea/snapshot_f4874ed7-fe85-
4d7d-b22b-082a2e068010/vm_id_9894f013-77dd-4514-8e65-818f4ae91d1f/
vm_res_id_9ae3a6e7-dffe-4424-badc-bc4de1a18b40_vda/a6289269-3e72-4085-
adca-e228ba656984
Format specific information:
  compat: 1.1
  lazy refcounts: false
  refcount bits: 16
  corrupt: false
```

As can be seen in the example, the backing file is an absolute path, which makes it necessary, that this path exists so the backing files can be accessed.

NetBackup for OpenStack is using the base64 hashing algorithm for the NFS mount-paths, to allow the configuration of multiple NFS Volumes at the same time. The hash value is calculated using the provided NFS path.

```
# echo -n 10.10.2.20:/upstream | base64  
MTAuMTAuMi4yMDovdXBzdHJlYW0=
```

If the path of the backing file is not available on the NetBackup for OpenStack VM and Compute nodes, the restores of incremental backups fails.

The tested and recommended method to make the backing files available is creating the required directory path and using `mount --bind` to make the path available for the backups.

```
#mount --bind <mount-path1> <mount-path2>
```

Running the `mount --bind` command will make the necessary path available until the next reboot. If it is required to have access to the path beyond a reboot, it is necessary to edit the `fstab`.

```
#vi /etc/fstab  
<mount-path1> <mount-path2> none bind 0 0
```

Example runbook for disaster recovery using NFS

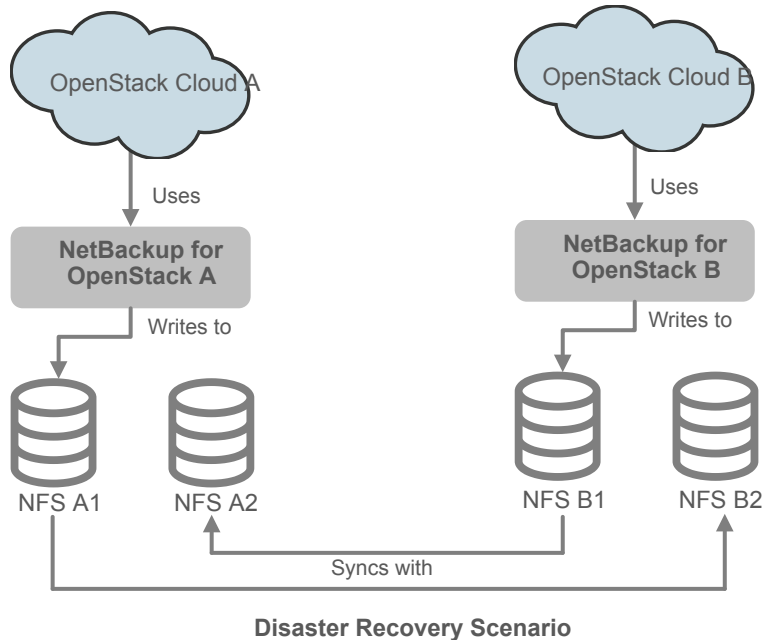
This runbook demonstrates how to set up disaster recovery with NetBackup for OpenStack for a given scenario.

Scenario

There are two OpenStack clouds available OpenStack Cloud A and OpenStack Cloud B". OpenStack Cloud B is the disaster recovery restore point of OpenStack Cloud A and vice versa. Both clouds have an independent NetBackup for OpenStack installation integrated. These NetBackup for OpenStack installations write their Backups to NFS targets. "NetBackup for OpenStack A" is writing to "NFS A1" and "NetBackup for OpenStack B" is writing to "NFS B1". The NFS Volumes used are getting synced against another NFS Volume on the other side. "NFS A1" is syncing with "NFS B2" and "NFS B1" is syncing with "NFS A2". The syncing process is set

up independently from NetBackup for OpenStack and will always favor the newer dataset.

Figure 7-2 Disaster recovery Scenario



This scenario covers the disaster recovery of a single policy and a complete Cloud. All processes are done by the OpenStack administrator.

Prerequisites for the disaster recovery process

This runbook assumes that the following is true:

- OpenStack Cloud A and OpenStack Cloud B both have an active NetBackup for OpenStack installation with a valid license
- OpenStack Cloud A and OpenStack Cloud B have free resources to host additional VMs
- OpenStack Cloud A and OpenStack Cloud B have Tenants/Projects available that are the designated restore points for Tenant/Projects of the other side
- Access to a user with the admin role permissions on domain level
- One of the OpenStack clouds is down/lost

We assume that OpenStack Cloud A is down and the policies are getting restored into OpenStack Cloud B.

In the case of the usage of shared Tenant networks, beyond the floating IP, the following additional requirement is needed: All Tenant Networks, Routers, Ports, Floating IPs, and DNS Zones are created.

Disaster recovery of a single policy

A single policy can do a disaster recovery in this Scenario, while both Clouds are still active. To do so the following high-level process needs to be followed:

1. Copy the policy directories to the configured NFS Volume.
2. Make the right Mount-Paths available.
3. Reassign the policy.
4. Restore the policy.
5. Clean up.

Copy the policy directories to the configured NFS Volume

Note: This process only shows how to get a policy from OpenStack Cloud A to OpenStack Cloud B. The vice versa process is similar.

As only a single policy is to be recovered it is more efficient to copy the data of that single policy over to the NFS B1 Volume, which is used by "NetBackup for OpenStack B".

Mount NFS B2 Volume to a NetBackup for OpenStack VM

It is recommended to use the NetBackup for OpenStack VM as a connector between both NFS Volumes, as the nova user is available on the NetBackup for OpenStack VM.

```
# mount <NFS B2-IP/NFS B2-FQDN>:<VOL-Path> /mnt
```

Identify the policy on the NFS B2 Volume

NetBackup for OpenStack are identified by their IDs and which they are stored on the Backup Target. See the following example:

```
workload_ac9cae9b-5e1b-4899-930c-6aa0600a2105
```

In case the policy ID is not known, it can be available in the backup metadata within the policy directories.

```

/.../workload_<id>/workload_db <<< Contains User ID and Project ID
of Workload owner
/.../workload_<id>/workload_vms_db <<< Contains VM IDs and VM Names
of all VMs actively protected be the Workload
    
```

Copy the policy

The identified policy needs to be copied with all subdirectories and files. Afterward, it is necessary to adjust the ownership to nova:nova with the right permissions.

```

# cp /mnt/policy_ac9cae9b-5e1b-4899-930c-6aa0600a2105 /var/
NetBackupOpenStack-mounts/MTAuMTAuMi4yMDovdXBzdHJlYW0=/
workload_ac9cae9b-5e1b-4899-930c-6aa0600a2105
# chown -R nova:nova /var/NetBackupOpenStack-mounts/
MTAuMTAuMi4yMDovdXBzdHJlYW0=/workload_ac9cae9b-5e1b-4899-
930c-6aa0600a2105
# chmod -R 644 /var/NetBackupOpenStack-mounts/
MTAuMTAuMi4yMDovdXBzdHJlYW0=/workload_ac9cae9b-5e1b-
4899-930c-6aa0600a2105
    
```

Make the Mount-Paths available

NetBackup for OpenStack backups are using qcow2 backing files, which make every incremental backup a full synthetic backup. These backing files can be made visible using the qemu-img tool.

```

#qemu-img info bd57ec9b-c4ac-4a37-a4fd-5c9aa002c778
image: bd57ec9b-c4ac-4a37-a4fd-5c9aa002c778
file format: qcow2
virtual size: 1.0G (1073741824 bytes)
disk size: 516K
cluster_size: 65536
    
```

```

backing file: /var/NetBackupOpenStack-mounts/MTAuMTAuMi4yMDovdXBzdHJlYW0=
/workload_ac9cae9b-5e1b-4899-930c-6aa0600a2105/snapshot_1415095d-
c047-400b-8b05-c88e57011263/vm_id_38b620f1-24ae-41d7-b0ab-85ffc2
d7958b/vm_res_id_d4ab3431-5ce3-4a8f-a90b-07606e2ffa33_vda/7c39
eb6a-6e42-418e-8690-b6368ecaa7bb
Format specific information:
    
```

```

compat: 1.1
lazy refcounts: false
refcount bits: 16
corrupt: false
    
```

The `MTAuMTAuMi4yMDovdXBzdHJlYW0=` part of the backing file path is the base64 hash value, which will be calculated upon the configuration of a NetBackup for OpenStack installation for each provided NFS-Share.

This hash value is calculated based on the provided NFS-Share path: `<NFS_IP>/<path>` If even one character in the NFS-Share path is different between the provided NFS-Share paths a completely different hash value is generated.

Policies that have moved between NFS-Shares require that their incremental backups can follow the same path as on their original Source Cloud. To achieve this it is necessary to create the mount path on all compute nodes of the Target Cloud.

Afterwards a mount bind is used to make the policy data accessible over the old and the new mount path. The following example shows the process of how to successfully identify the necessary mount points and create the mount bind.

Identify the base64 hash values

The used hash values can be calculated using the base64 tool in any Linux distribution.

```
# echo -n 10.10.2.20:/NFS_A1 | base64
MTAuMTAuMi4yMDovdXBzdHJlYW1fc291cmNl
    
```

```
# echo -n 10.20.3.22:/NFS_B2 | base64
MTAuMjAuMy4yMjovdXBzdHJlYW1fdGFyZ2V0
    
```

Create and bind the paths

Based on the identified base64 hash values the following paths are required on each Compute node.

```
/var/NetBackupOpenStack-mounts/MTAuMTAuMi4yMDovdXBzdHJlYW1fc291cmNl
```

and

```
/var/NetBackupOpenStack-mounts/MTAuMjAuMy4yMjovdXBzdHJlYW1fdGFyZ2V0
```

In the scenario the mount path of the NFS Share_A1 needs to be created and bound to the target cloud.

```
#mkdir /var/NetBackupOpenStack-mounts/  
MTAuMTAuMi4yMDovdXBzdHJlYW1fc291cmNl  
#mount --bind  
/var/NetBackupOpenStack-mounts/MTAuMjAuMy4yMjovdXBzdHJlYW1fdGFyZ2V0/  
/var/NetBackupOpenStack-mounts/MTAuMTAuMi4yMDovdXBzdHJlYW1fc291cmNl
```

To keep the desired mount past a restart it is recommended to edit the `fstab` of all compute nodes accordingly.

```
#vi /etc/fstab  
/var/NetBackupOpenStack-mounts/MTAuMjAuMy4yMjovdXBzdHJlYW1fdGFyZ2V0/  
/var/NetBackup for OpenStack-mounts/ MTAuMTAuMi4yMDovdXBzdHJlYW1fc291cmNl  
none          bind          0 0
```

Reassign the policy

NetBackup for OpenStack policies have clear ownership. When a policy is moved to a different cloud it is necessary to change the ownership. The ownership can only be changed by OpenStack administrators.

Add admin-user to required domains and projects

To fulfill the required tasks an admin role user is used. This user will be used until the policy has been restored. Therefore, it is necessary to provide this user access to the desired Target Project on the Target Cloud.

```
# source {customer admin rc file}  
# OpenStack role add Admin --user <my_admin_user> --user-domain  
<admin_domain> --domain <target_domain>  
# OpenStack role add Admin --user <my_admin_user> --user-domain  
<admin_domain> --project <target_project> --project-domain <target_domain>  
# OpenStack role add <Backup Trustee Role> --user <my_admin_user>  
--user-domain <admin_domain> --project <destination_project>  
--project-domain <target_domain>
```

Discover orphaned policies from NFS-Storage of Target Cloud

Each NetBackup for OpenStack installation maintains a database of policies that are known to the NetBackup for OpenStack installation. Policies that are not maintained by a specific NetBackup for OpenStack installation, are from the perspective of that installation, orphaned policies. An orphaned policy is a policy accessible on the NFS-Share, that is not assigned to any existing project in the Cloud the NetBackup for OpenStack installation is protecting.

```
# workloadmgr workload-get-orphaned-workloads-list --migrate_cloud True
```

List available projects on Target Cloud in the Target Domain

The identified orphaned policies need to be assigned to their new projects. The following provides the list of all available projects viewable by the used admin-user in the target_domain.

```
# OpenStack project list --domain <target_domain>
```

List available users on the Target Cloud in the Target Project that have the right backup trustee role

To allow project owners to work with the policies and ensure that the policy is assigned to the user with the backup trustee role.

```
# OpenStack role assignment list --project <target_project>  
--project-domain <target_domain> --role <backup_trustee_role>
```

Reassign the policy to the target project

Now that all information is gathered the policy can be reassigned to the target project.

```
# workloadmgr workload-reassign-workloads --new_tenant_id  
{target_project_id} --user_id {target_user_id} --workload_ids  
{workload_id} --migrate_cloud True
```

Verify that the policy is available at the desired target_project

After the policy has been assigned to the new project it is recommended to verify that the policy is managed by the Target NetBackup for OpenStack and is assigned to the right project and user.

```
# workloadmgr workload-show ac9cae9b-5e1b-4899-930c-6aa0600a2105
```

Restore the policy

The reassigned policy can be restored using Horizon following the Selective Restore procedure. See [“Selective Restore”](#) on page 88.

This runbook will continue on the CLI only path.

Prepare the selective restore by getting the snapshot information

To be able to do the necessary selective restore a few pieces of information about the snapshot to be restored are required. The following process will provide all necessary information.

List all Snapshots of the policy to restore to identify the snapshot to restore

```
# workloadmgr snapshot-list --workload_id ac9cae9b-5e1b-4899-930c-6aa0600a2105 --all True
```

Get Snapshot Details with network details for the desired snapshot

```
# workloadmgr snapshot-show --output networks 7e39e544-537d-4417-853d-11463e7396f9
```

Get Snapshot Details with disk details for the desired Snapshot.

```
[root@upstreamcontroller ~(keystone_admin)]# workloadmgr snapshot-show --output disks 7e39e544-537d-4417-853d-11463e7396f9
```

Prepare the selective restore by creating the restore.json file

The selective restore is using a restore.json file for the CLI command. This restore.json file needs to be adjusted according to the desired restore.

```
{
  u'description':u'<description of the restore>',
  u'oneclickrestore':False,
  u'restore_type':u'selective',
  u'type':u'OpenStack',
  u'name':u'<name of the restore>'
  u'OpenStack':{
    u'instances':[
      {
        u'name':u'<name instance 1>',
        u'availability_zone':u'<AZ instance 1>',
        u'nics':[ #####Leave empty for network topology restore
        ],
        u'vdisks':[
          {
            u'id':u'<old disk id>',
            u'new_volume_type':u'<new volume type name>',
            u'availability_zone':u'<new cinder volume AZ>'
          }
        ]
      }
    ]
  }
}
```

```

    }
  ],
  u'flavor':{
    u'ram':<RAM in MB>,
    u'ephemeral':<GB of ephemeral disk>,
    u'vcpus':<# vCPUs>,
    u'swap':u'<GB of Swap disk>',
    u'disk':<GB of boot disk>,
    u'id':u'<id of the flavor to use>'
  },
  u'include':<True/False>,
  u'id':u'<old id of the instance>'
} #####Repeat for each instance in the snapshot
],
u'restore_topology':<True/False>,
u'networks_mapping':{
  u'networks':[ #####Leave empty for network topology restore

]
}
}
}
}
}

```

Run the selective restore

To do the actual restore use the following command:

```
# workloadmgr snapshot-selective-restore --filename restore.json
{snapshot id}
```

Verify the restore

To verify the success of the restore from a NetBackup for OpenStack perspective the restore status is checked.

```
[root@upstreamcontroller ~(keystone_admin)]# workloadmgr restore-list
--snapshot_id 5928554d-a882-4881-9a5c-90e834c071af
```

Clean up

After the Disaster Recovery Process has been successfully completed it is recommended to bring the nbosvm installation back into its original state to be ready for the next DR process.

Delete the policy

Delete the policy that got restored.

```
# workloadmgr workload-delete <workload_id>
```

Remove the database entry

The NetBackup for OpenStack database is following the OpenStack standard of not deleting any database entries upon deletion of the cloud object. Any policy, Snapshot or Restore, which gets deleted, is marked as deleted only.

To allow the NetBackup for OpenStack installation to be ready for another disaster recovery it is necessary to completely delete the entries of the policies, which have been restored.

NetBackup for OpenStack does provide and maintain a script to safely delete policy entries and all connected entities from the NetBackup for OpenStack database.

This script can be found here:

<https://github.com/nbosData/solutions/tree/master/OpenStack/CleanWlmDatabase>

Remove the admin user from the project

After all restores for the target project have been achieved it is recommended to remove the used admin user from the project again.

```
# source {customer admin rc file}
# OpenStack role remove Admin --user <my_admin_user> --user-domain
<admin_domain> --domain <target_domain>
# OpenStack role remove Admin --user <my_admin_user> --user-domain
<admin_domain> --project <target_project> --project-domain <target_domain>
# OpenStack role remove <Backup Trustee Role> --user <my_admin_user>
--user-domain <admin_domain> --project <destination_project>
--project-domain <target_domain>
```

Disaster recovery of a complete cloud

This Scenario will cover the disaster recovery of a full cloud. It is assumed that the source cloud is down or lost completely. To do the disaster recovery the following high-level process needs to be followed:

1. Reconfigure the Target NetBackup for OpenStack installation.
2. Make the right Mount-Paths available.
3. Reassign the policy.

4. Restore the policy.
5. Reconfigure the Target NetBackup for OpenStack installation back to the original one.
6. Clean up.

Reconfigure the Target NetBackup for OpenStack installation

Before the disaster recovery Process can start it is necessary to make the backups to be restored available for the NetBackup for OpenStack installation. The following steps need to be done to completely reconfigure the NetBackup for OpenStack installation.

During the reconfiguration process all backups of the Target Region will be on hold and it is not recommended to create new backup jobs until the disaster recovery Process has finished and the original NetBackup for OpenStack configuration has been restored.

Add NFS B2 to the NetBackup for OpenStack Appliance Cluster

To add the NFS-B2 to the NetBackup for OpenStack Appliance cluster the NetBackup for OpenStack can either be fully reconfigured to use both NFS Volumes or it is possible to edit the configuration file and then restart all services. This procedure describes how to edit the conf file and restart the services. This needs to be repeated on every NetBackup for OpenStack Appliance.

Edit the workloadmgr.conf

```
# vi /etc/workloadmgr/workloadmgr.conf
```

Look for the line defining the NFS mounts

```
vault_storage_nfs_export = <NFS_B1/NFS_B1-FQDN>:/<VOL-B1-Path>
```

Add NFS B2 to that as comma-separated list. Space is not necessary, but can be set.

```
vault_storage_nfs_export = <NFS-IP/NFS-FQDN>:/<VOL-1-Path>,  
<NFS-IP/NFS-FQDN>:/<VOL-2-Path>
```

Write and close the workloadmgr.conf

Restart the wlm-workloads service

```
# systemctl restart wlm-workloads
```

Add NFS B2 to the NetBackup for OpenStack Datamovers

NetBackup for OpenStack integrates natively into the OpenStack deployment tools. When using the Red Hat director, it is recommended to adapt the environment files for these orchestrators and update the Datamovers through them.

To add the NFS B2 to the NetBackup for OpenStack Datamovers manually the `nbosdm.conf` file needs to be edited and the service restarted.

Edit the `nbosdm.conf`.

```
# vi /etc/nbosdm/nbosdm.conf
```

Look for the line defining the NFS mounts.

```
vault_storage_nfs_export = <NFS_B1-IP/NFS_B1-FQDN>:./<VOL-B1-Path>
```

Add NFS B2 to that as comma-separated list. Space is not necessary, but can be set.

```
vault_storage_nfs_export = <NFS_B1-IP/NFS-FQDN>:./<VOL-B1-Path>,  
<NFS_B2-IP/NFS-FQDN>:./<VOL-B2-Path>
```

Write and close the `nbosdm.conf`

Restart the `nbosdm` service.

```
# systemctl restart nbosdm
```

Make the Mount-Paths available

NetBackup for OpenStack backups are using `qcow2` backing files, which make every incremental backup a full synthetic backup. These backing files can be made visible using the `qemu-img` tool.

```
#qemu-img info bd57ec9b-c4ac-4a37-a4fd-5c9aa002c778  
image: bd57ec9b-c4ac-4a37-a4fd-5c9aa002c778  
file format: qcow2  
virtual size: 1.0G (1073741824 bytes)  
disk size: 516K
```

```
cluster_size: 65536

backing file: /var/NetBackupOpenStack-mounts/MTAuMTAuMi4yMDovdXBzdHJlYW0=
/workload_ac9cae9b-5e1b-4899-930c-6aa0600a2105/snapshot_1415095d
-c047-400b-8b05-c88e57011263/vm_id_38b620f1-24ae-41d7-b0ab-85ffc
2d7958b/vm_res_id_d4ab3431-5ce3-4a8f-a90b-07606e2ffa33_vda/7c39eb
6a-6e42-418e-8690-b6368ecaa7bb
Format specific information:
    compat: 1.1
    lazy refcounts: false
    refcount bits: 16
    corrupt: false
```

The MTAuMTAuMi4yMDovdXBzdHJlYW0= part of the backing file path is the base64 hash value, which will be calculated upon the configuration of a NetBackup for OpenStack installation for each provided NFS-Share.

This hash value is calculated based on the provided NFS-Share path:
 <NFS_IP>/<path> If even one character in the NFS-Share path is different between the provided NFS-Share paths a completely different hash value is generated.

Policies that have moved between NFS-Shares require that their incremental backups can follow the same path as on their original Source Cloud. To achieve this it is necessary to create the mount path on all compute nodes of the Target Cloud.

Afterwards a mount bind is used to make the policies data accessible over the old and the new mount path. The following example shows the process of how to successfully identify the necessary mount points and create the mount bind.

Identify the base64 hash values

The used hash values can be calculated using the base64 tool in any Linux distribution.

```
# echo -n 10.10.2.20:/NFS_A1 | base64
MTAuMTAuMi4yMDovdXBzdHJlYW1fc291cmNl

# echo -n 10.20.3.22:/NFS_B2 | base64
MTAuMjAuMy4yMjovdXBzdHJlYW1fdGFyZ2V0
```

Create and bind the paths

Based on the identified base64 hash values the following paths are required on each Compute node.

```
/var/NetBackupOpenStack-mounts/MTAuMTAuMi4yMDovdXBzdHJlYW1fc291cmNl
```

and

```
/var/NetBackupOpenStack-mounts/MTAuMjAuMy4yMjovdXBzdHJlYW1fdGFyZ2V0
```

In the scenario the mount path of the NFS Share_A1 needs to be created and bound to the target cloud.

```
#mkdir /var/NetBackupOpenStack-mounts/MTAuMTAuMi4yMDovdXBzdHJlYW1fc291cmNl
#mount --bind
/var/NetBackupOpenStack-mounts/MTAuMjAuMy4yMjovdXBzdHJlYW1fdGFyZ2V0/
/var/NetBackup for OpenStack-mounts/MTAuMTAuMi4yMDovdXBzdHJlYW1fc291cmNl
```

To keep the desired mount past a restart it is recommended to edit the fstab of all compute nodes accordingly.

```
#vi /etc/fstab
/var/NetBackupOpenStack-mounts/MTAuMjAuMy4yMjovdXBzdHJlYW1fdGFyZ2V0/
/ var/NetBackup for OpenStack-mounts/ MTAuMTAuMi4yMDovdXBzdHJlYW1fc291cmNl
none          bind          0 0
```

Reassign the policy

NetBackup for OpenStack policies have clear ownership. When a policy is moved to a different cloud it is necessary to change the ownership. The ownership can only be changed by OpenStack administrators.

Add admin-user to required domains and projects

To fulfill the required tasks an admin role user is used. This user is used until the policy is restored. Therefore, it is necessary to provide this user access to the desired Target Project on the Target Cloud.

```
# source {customer admin rc file}
# OpenStack role add Admin --user <my_admin_user> --user-domain
<admin_domain> --domain <target_domain>
# OpenStack role add Admin --user <my_admin_user> --user-domain
<admin_domain> --project <target_project> --project-domain <target_domain>
# OpenStack role add <Backup Trustee Role> --user <my_admin_user>
--user-domain <admin_domain> --project <destination_project>
--project-domain <target_domain>
```

Discover orphaned policies from NFS-Storage of Target Cloud

Each NetBackup for OpenStack installation maintains a database of policies that are known to the NetBackup for OpenStack installation. Policies that are not maintained by a specific NetBackup for OpenStack installation, are from the perspective of that installation, orphaned policies. An orphaned policy is a policy accessible on the NFS-Share, that is not assigned to any existing project in the Cloud the NetBackup for OpenStack installation is protecting.

```
# workloadmgr workload-get-orphaned-workloads-list --migrate_cloud True
```

List available projects on Target Cloud in the Target Domain

The identified orphaned policies need to be assigned to their new projects. The following provides the list of all available projects viewable by the used admin-user in the target_domain.

```
# OpenStack project list --domain <target_domain>
```

List available users on the Target Cloud in the Target Project that have the right backup trustee role

To allow project owners to work with the policies and ensure that the policy is assigned to the user with the backup trustee role.

```
# OpenStack role assignment list --project <target_project>  
--project-domain <target_domain> --role <backup_trustee_role>
```

Reassign the policy to the target project

Now that all information is gathered the policy can be reassigned to the target project.

```
# workloadmgr workload-reassign-workloads --new_tenant_id  
{target_project_id} --user_id {target_user_id} --workload_ids  
{workload_id} --migrate_cloud True
```

Verify that the policy is available at the desired target_project

After the policy has been assigned to the new project it is recommended to verify that the policy is managed by the Target NetBackup for OpenStack and is assigned to the right project and user.

```
# workloadmgr workload-show ac9cae9b-5e1b-4899-930c-6aa0600a2105
```

Restore the policy

The reassigned policy can be restored using Horizon following the Selective Restore procedure.

This runbook will continue on the CLI only path.

Prepare the selective restore by getting the snapshot information

To be able to do the necessary selective restore a few pieces of information about the snapshot to be restored are required. The following process will provide all necessary information.

List all Snapshots of the policy to restore to identify the snapshot to restore.

```
# workloadmgr snapshot-list --workload_id ac9cae9b-5e1b-4899-930c-6aa0600a2105 --all True
```

Get Snapshot Details with network details for the desired snapshot.

```
# workloadmgr snapshot-show --output networks 7e39e544-537d-4417-853d-11463e7396f9
```

Get Snapshot Details with disk details for the desired Snapshot.

```
[root@upstreamcontroller ~(keystone_admin)]# workloadmgr snapshot-show --output disks 7e39e544-537d-4417-853d-11463e7396f9
```

Prepare the selective restore by creating the restore.json file

The selective restore is using a restore.json file for the CLI command. This restore.json file needs to be adjusted according to the desired restore.

```
{
  u'description':u'<description of the restore>',
  u'oneclickrestore':False,
  u'restore_type':u'selective',
  u'type':u'OpenStack',
  u'name':u'<name of the restore>'
  u'OpenStack':{
    u'instances':[
      {
        u'name':u'<name instance 1>',
        u'availability_zone':u'<AZ instance 1>',
        u'nics':[ #####Leave empty for network topology restore
        ],
        u'vdisks':[
          {
            u'id':u'<old disk id>',
            u'new_volume_type':u'<new volume type name>',
```


Reconfigure the Target NetBackup for OpenStack installation back to the original one

After the Disaster Recovery Process has finished it is necessary to return the NetBackup for OpenStack installation to its original configuration. The following steps need to be done to completely reconfigure the NetBackup for OpenStack installation.

During the reconfiguration process will all backups of the Target Region be on hold and it is not recommended to create new backup jobs until the Disaster Recovery Process has finished and the original NetBackup for OpenStack configuration has been restored.

Delete NFS B2 from the NetBackup for OpenStack Appliance Cluster

To delete the NFS-B2 to the NetBackup for OpenStack Appliance cluster the NetBackup for OpenStack can either be fully reconfigured to use both NFS Volumes or it is possible to edit the configuration file and then restart all services. This procedure describes how to edit the conf file and restart the services. This needs to be repeated on every NetBackup for OpenStack Appliance.

See [“Configuring NetBackup for OpenStack”](#) on page 38.

Edit the `workloadmgr.conf`.

```
# vi /etc/workloadmgr/workloadmgr.conf
```

Look for the line defining the NFS mounts.

```
vault_storage_nfs_export = <NFS_B1-IP/NFS-FQDN>:/<VOL-B1-Path>,  
<NFS_B2-IP/NFS-FQDN>:/<VOL-B2-Path>
```

Delete NFS B2 from the comma-separated list.

```
vault_storage_nfs_export = <NFS_B1-IP/NFS_B1-FQDN>:/<VOL-B1-Path>
```

Write and close the `workloadmgr.conf`.

Restart the `wlm-workloads` service.

```
# systemctl restart wlm-workloads
```

Delete NFS B2 from the NetBackup for OpenStack Datamovers

Warning: NetBackup for OpenStack is integrating natively into the OpenStack deployment tools. When using the Red Hat director, it is recommended to adapt the environment files for these orchestrators and update the Datamovers through them.

To delete the NFS B2 to the NetBackup for OpenStack Datamovers manually the `nbosdm.conf` file needs to be edited and the service restarted.

Edit the `nbosdm.conf`.

```
# vi /etc/nbosdm/nbosdm.conf
```

Look for the line defining the NFS mounts.

```
vault_storage_nfs_export = <NFS_B1-IP/NFS-FQDN>:/<VOL-B1-Path>,<NFS_B2-IP/NFS
```

Delete NFS B2 from the comma-separated list.

```
vault_storage_nfs_export = <NFS-IP/NFS-FQDN>:/<VOL-1-Path>
```

Write and close the `nbosdm.conf`.

Restart the `nbosdm` service.

```
# systemctl restart nbosdm
```

Clean up

After the disaster recovery Process has been successfully completed and the NetBackup for OpenStack installation that is reconfigured to its original state, it is recommended to do the following additional steps to be ready for the next disaster recovery process.

Remove the database entry

The NetBackup for OpenStack database is following the OpenStack standard of not deleting any database entries upon deletion of the cloud object. Any policy, Snapshot or Restore, which gets deleted, is marked as deleted only.

To allow the NetBackup for OpenStack installation to be ready for another disaster recovery it is necessary to completely delete the entries of the policies, which have been restored.

NetBackup for OpenStack does provide and maintain a script to safely delete policy entries and all connected entities from the NetBackup for OpenStack database.

This script can be found here:

<https://github.com/nbosData/solutions/tree/master/OpenStack/CleanWlmDatabase>

Remove the admin user from the project

After all restores for the target project have been achieved it is recommended to remove the used admin user from the project again.

```
# source {customer admin rc file}
# OpenStack role remove Admin --user <my_admin_user> --user-domain
<admin_domain> --domain <target_domain>
# OpenStack role remove Admin --user <my_admin_user> --user-domain
<admin_domain> --project <target_project> --project-domain <target_domain>
# OpenStack role remove <Backup Trustee Role> --user <my_admin_user>
--user-domain <admin_domain> --project <destination_project>
--project-domain <target_domain>
```

Troubleshooting

This chapter includes the following topics:

- [General Troubleshooting Tips](#)
- [Using the workloadmgr CLI tool on the NetBackup for OpenStack Appliance](#)
- [Health check of NetBackup for OpenStack](#)
- [Important log files](#)

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 policy-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 gather the metadata about the backed-up VMs 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 VM 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 datamover 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 dmapi service connects to the nova api database requesting the compute host of a provided VM.

Once the compute host has been identified the nbosdmapi forwards the command from the NetBackup for OpenStack Cluster to the datamover running on the identified compute host.

nbosdm

The datamover is the NetBackup for OpenStack service running on the compute nodes.

Each datamover is responsible for the VMs running on top of its compute node. A datamover cannot work with VMs running on a different compute node.

The datamover controls the freeze and thaw of VMs as well as the actual movement of the data.

Everything on the Backup Target happens as 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 backup 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 policies, backups, or restores.

OpenStack Quotas

NetBackup for OpenStack creates Cinder Snapshots and temporary Cinder Volumes. The OpenStack Quotas need to allow that.

Every disk that is backed up requires one temporary Cinder Volume.

Every Cinder Volume that is backed up requires two Cinder Snapshots. The second Cinder Snapshot is temporary to calculate the incremental.

Using the workloadmgr CLI tool on the NetBackup for OpenStack Appliance

To use the workloadmgr CLI tool on the NetBackup for OpenStack appliance it is only necessary to activate the virtual environment of the workloadmgr

```
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

wlm-workloads

This service runs and is active on every NetBackup for OpenStack node.

```
[root@Upstream ~]# systemctl status wlm-workloads
● wlm-workloads.service - workloadmanager workloads service
   Loaded: loaded (/etc/systemd/system/wlm-workloads.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 (workloadmgr-wor)
    Tasks: 17
   CGroup: /system.slice/wlm-workloads.service
           └─12779 /home/stack/myansible/bin/python
/home/stack/myansible/bin/workloadmgr-workloads
--config-file=/etc/workloadmgr/workloadmgr.conf
           └─12982 /home/stack/myansible/bin/python
/home/stack/myansible/bin/workloadmgr-workloads
--config-file=/etc/workloadmgr/workloadmgr.conf
           └─12983 /home/stack/myansible/bin/python
/home/stack/myansible/bin/workloadmgr-workloads
--config-file=/etc/workloadmgr/workloadmgr.conf
           └─12984 /home/stack/myansible/bin/python
/home/stack/myansible/bin/workloadmgr-workloads
--config-file=/etc/workloadmgr/workloadmgr.conf
           [...]

```

wlm-api

This service runs and is active on every NetBackup for OpenStack node.

```
[root@Upstream ~]# systemctl status wlm-api
● wlm-api.service - Cluster Controlled wlm-api
   Loaded: loaded (/etc/systemd/system/wlm-api.service; disabled;
   vendor preset: disabled)
   Drop-In: /run/systemd/system/wlm-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 (workloadmgr-api)

```

```

Tasks: 1
CGroup: /system.slice/wlm-api.service
└─11815 /home/stack/myansible/bin/python /home/stack/
myansible/bin/workloadmgr-api --config-file=/etc/
workloadmgr/workloadmgr.conf

```

wlm-scheduler

This service runs and is active on every NetBackup for OpenStack node.

```

[root@Upstream ~]# systemctl status wlm-scheduler
● wlm-scheduler.service - Cluster Controlled wlm-scheduler
   Loaded: loaded (/etc/systemd/system/wlm-scheduler.service; disabled;
   vendor preset: disabled)
   Drop-In: /run/systemd/system/wlm-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 (workloadmgr-sch)
     Tasks: 1
    CGroup: /system.slice/wlm-scheduler.service
            └─29439 /home/stack/myansible/bin/python /home/stack/myansible
            /bin/workloadmgr-scheduler --config-file=/etc/workloadmgr/
            workloadmgr.conf

```

wlm-cron

This service is controlled by pacemaker and runs only on the master node

```

[root@Upstream ~]# systemctl status wlm-cron
● wlm-cron.service - Cluster Controlled wlm-cron
   Loaded: loaded (/etc/systemd/system/wlm-cron.service; disabled;
   vendor preset: disabled)
   Drop-In: /run/systemd/system/wlm-cron.service.d
            └─50-pacemaker.conf
   Active: active (running) since Wed 2021-01-27 19:59:26 UTC; 6 days ago
   Main PID: 23071 (workloadmgr-cro)
     CGroup: /system.slice/wlm-cron.service
            └─23071 /home/stack/myansible/bin/python3 /home/stack/
myansible/bin/workloadmgr-cron --config-file=/etc/workloadmgr
/workloadmgr.conf
            └─23248 /home/stack/myansible/bin/python3 /home/stack/
myansible/bin/workloadmgr-cron --config-file=/etc/workloadmgr/

```

```
workloadmgr.conf
```

```
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant workloadmgr-cron
[23071]: ● wlm-cron.service - Cluster Controlled wlm-cron
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant workloadmgr-cron
[23071]: Loaded: loaded (/etc/systemd/system/wlm-cron.service; disabled;
vendor preset: disabled)
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant workloadmgr-cron
[23071]: Drop-In: /run/systemd/system/wlm-cron.service.d
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant workloadmgr-cron
[23071]: └─50-pacemaker.conf
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant workloadmgr-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 workloadmgr-cron
[23071]: Main PID: 23071 (workloadmgr-cro)
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant workloadmgr-cron
[23071]: CGroup: /system.slice/wlm-cron.service
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant workloadmgr-cron
[23071]: └─23071 /home/stack/myansible/bin/python3 /home/stack/myansible/
bin/workloadmgr-cron --config-file=/etc/workloadmgr/workloadmgr.conf
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant workloadmgr-cron
[23071]: └─23248 /home/stack/myansible/bin/python3 /home/stack/myansible/
bin/workloadmgr-cron --config-file=/etc/workloadmgr/workloadmgr.conf
Feb 03 19:28:43 nbosvml-ansible-ussuri-ubuntu18-vagrant workloadmgr-cron
[23071]: └─27145 /usr/bin/systemctl status wlm-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 wlm-api and wlm-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: nbosvml-ansible-ussuri-ubuntu18-vagrant (version
1.1.23-1.el7_9.1-9acf116022) - chaptererition 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
wlm-cron        (systemd:wlm-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
| e9aa6630bf74a9bb7562d4161f4e07d | 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/nbosdmapi.service
           └─11252 /usr/bin/python /usr/bin/nbosdmapi-api
           └─11280 /usr/bin/python /usr/bin/nbosdmapi-api

```

The datamover service

The datamover 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 - nbosdm container
   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 `workloadmgr-workloads.log`, which contains all logs about ongoing and past NetBackup for OpenStack backup and restore tasks. It can be found at:

```
/var/log/workloadmgr/nbosjm-workloads.log
```

The next important log is the `workloadmgr-api.log`, which contains all logs about API calls received by the NetBackup for OpenStack Cluster. It can be found at:

```
/var/log/workloadmgr/nbosjm-api.log
```

The log for the third service is the `workloadmgr-scheduler.log`, which contains all logs about the internal job scheduling between NetBackup for OpenStack nodes in the NetBackup for OpenStack Cluster.

```
/var/log/workloadmgr/nbosjm-scheduler.log
```

The last service running on the NetBackup for OpenStack Nodes is the `wlm-cron` service, which controls the scheduled automated backups.

```
/var/log/workloadmgr/nbosjm-workloads.log
```

In case of using S3 as a backup target, there is also a log file that keeps track of the S3-Fuse plug-in that is used to connect with the S3 storage.

```
/var/log/workloadmgr/s3vaultfuse.py.log
```

NetBackup for OpenStack Datamover service logs on RHOSP

- **Datamover API log**
 The log for the NetBackup for OpenStack Datamover API service is located on the nodes, typically controller, where the NetBackup for OpenStack Datamover API container is running under:

```
/var/log/containers/nbosdmapi/nbosdmapi.log
```
- **Datamover log**
 The log for the NetBackup for OpenStack Datamover service is located on the nodes, typically compute, where the NetBackup for OpenStack Datamover container is running under:

```
/var/log/containers/nbosdm/nbosdm.log
```

 In case of S3 being used in the log for the S3 Fuse plug-in that is located on the same nodes under:

```
/var/log/containers/nbosdm/nbos-object-store.log
```

NetBackup for OpenStack Datamover service logs on Ansible OpenStack

- **Datamover API log**
 The log for the NetBackup for OpenStack Datamover API service is located on the nodes, typically controller, where the NetBackup for OpenStack Datamover API container is running. Log into the nbosdmapi container using `lxc-attach` command.

```
lxc-attach -n controller_nbosdmapi_container-a11984bf
```

 The log file is then located under:

```
/var/log/nbosdmapi/nbosdmapi.log
```
- **Datamover log**
 The log for the NetBackup for OpenStack Datamover service is typically located on the compute nodes and the logs can be found here:

```
/var/log/nbosdm/nbosdm.log
```

 In case of S3 being used in the log for the S3 Fuse plug-in that is located on the same nodes under:

```
/var/log/nbos-object-store/nbos-object-store.log
```