



Red Hat OpenStack Platform 16.1

Release Notes

Release details for Red Hat OpenStack Platform 16.1

Red Hat OpenStack Platform 16.1 Release Notes

Release details for Red Hat OpenStack Platform 16.1

OpenStack Documentation Team
Red Hat Customer Content Services
rhos-docs@redhat.com

Legal Notice

Copyright © 2020 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This document outlines the major features, enhancements, and known issues in this release of Red Hat OpenStack Platform.

Table of Contents

CHAPTER 1. INTRODUCTION	3
1.1. ABOUT THIS RELEASE	3
1.2. REQUIREMENTS	3
1.3. DEPLOYMENT LIMITS	3
1.4. DATABASE SIZE MANAGEMENT	4
1.5. CERTIFIED DRIVERS AND PLUG-INS	4
1.6. CERTIFIED GUEST OPERATING SYSTEMS	4
1.7. PRODUCT CERTIFICATION CATALOG	4
1.8. BARE METAL PROVISIONING OPERATING SYSTEMS	4
1.9. HYPERVISOR SUPPORT	4
1.10. CONTENT DELIVERY NETWORK (CDN) REPOSITORIES	4
1.10.1. Undercloud repositories	5
1.10.2. Overcloud repositories	7
1.11. PRODUCT SUPPORT	10
CHAPTER 2. TOP NEW FEATURES	11
2.1. COMPUTE	11
2.2. DISTRIBUTED COMPUTE NODES (DCN)	11
2.3. NETWORKING	11
2.4. STORAGE	11
2.5. BARE METAL SERVICE	12
2.6. NETWORK FUNCTIONS VIRTUALIZATION	12
2.7. TECHNOLOGY PREVIEWS	12
CHAPTER 3. RELEASE INFORMATION	14
3.1. RED HAT OPENSTACK PLATFORM 16.1 GA	14
3.1.1. Bug Fix	14
3.1.2. Enhancements	16
3.1.3. Technology Preview	18
3.1.4. Rebase: Bug Fixes and Enhancements	20
3.1.5. Release Notes	20
3.1.6. Known Issues	23
3.1.7. Removed Functionality	28
CHAPTER 4. TECHNICAL NOTES	29
4.1. RHEA-2020:3148 RED HAT OPENSTACK PLATFORM 16.1 GENERAL AVAILABILITY ADVISORY	29

CHAPTER 1. INTRODUCTION

1.1. ABOUT THIS RELEASE

This release of Red Hat OpenStack Platform is based on the OpenStack "Train" release. It includes additional features, known issues, and resolved issues specific to Red Hat OpenStack Platform.

Only changes specific to Red Hat OpenStack Platform are included in this document. The release notes for the OpenStack "Train" release itself are available at the following location:

<https://releases.openstack.org/train/index.html>.

Red Hat OpenStack Platform uses components from other Red Hat products. For specific information pertaining to the support of these components, see

<https://access.redhat.com/site/support/policy/updates/openstack/platform/>.

To evaluate Red Hat OpenStack Platform, sign up at <http://www.redhat.com/openstack/>.



NOTE

The Red Hat Enterprise Linux High Availability Add-On is available for Red Hat OpenStack Platform use cases. For more details about the add-on, see <http://www.redhat.com/products/enterprise-linux-add-ons/high-availability/>. For details about the package versions to use in combination with Red Hat OpenStack Platform, see <https://access.redhat.com/site/solutions/509783>.

1.2. REQUIREMENTS

This version of Red Hat OpenStack Platform runs on the most recent fully supported release of Red Hat Enterprise Linux 8.2.

The Red Hat OpenStack Platform dashboard is a web-based interface that allows you to manage OpenStack resources and services.

The dashboard for this release supports the latest stable versions of the following web browsers:

- Chrome
- Mozilla Firefox
- Mozilla Firefox ESR
- Internet Explorer 11 and later (with **Compatibility Mode** disabled)



NOTE

Prior to deploying Red Hat OpenStack Platform, it is important to consider the characteristics of the available deployment methods. For more information, see [Installing and Managing Red Hat OpenStack Platform](#).

1.3. DEPLOYMENT LIMITS

For a list of deployment limits for Red Hat OpenStack Platform, see [Deployment Limits for Red Hat OpenStack Platform](#).

1.4. DATABASE SIZE MANAGEMENT

For recommended practices on maintaining the size of the MariaDB databases in your Red Hat OpenStack Platform environment, see [Database Size Management for Red Hat Enterprise Linux OpenStack Platform](#).

1.5. CERTIFIED DRIVERS AND PLUG-INS

For a list of the certified drivers and plug-ins in Red Hat OpenStack Platform, see [Component, Plug-In, and Driver Support in Red Hat OpenStack Platform](#).

1.6. CERTIFIED GUEST OPERATING SYSTEMS

For a list of the certified guest operating systems in Red Hat OpenStack Platform, see [Certified Guest Operating Systems in Red Hat OpenStack Platform and Red Hat Enterprise Virtualization](#).

1.7. PRODUCT CERTIFICATION CATALOG

For a list of the Red Hat Official Product Certification Catalog, see [Product Certification Catalog](#).

1.8. BARE METAL PROVISIONING OPERATING SYSTEMS

For a list of the guest operating systems that can be installed on bare metal nodes in Red Hat OpenStack Platform through Bare Metal Provisioning (ironic), see [Supported Operating Systems Deployable With Bare Metal Provisioning \(ironic\)](#).

1.9. HYPERVISOR SUPPORT

This release of the Red Hat OpenStack Platform is supported only with the **libvirt** driver (using KVM as the hypervisor on Compute nodes).

This release of the Red Hat OpenStack Platform runs with Bare Metal Provisioning.

Bare Metal Provisioning has been fully supported since the release of Red Hat OpenStack Platform 7 (Kilo). Bare Metal Provisioning allows you to provision bare-metal machines using common technologies (such as PXE and IPMI) to cover a wide range of hardware while supporting pluggable drivers to allow the addition of vendor-specific functionality.

Red Hat does not provide support for other Compute virtualization drivers such as the deprecated VMware "direct-to-ESX" hypervisor or non-KVM libvirt hypervisors.

1.10. CONTENT DELIVERY NETWORK (CDN) REPOSITORIES

This section describes the repositories required to deploy Red Hat OpenStack Platform 16.1.

You can install Red Hat OpenStack Platform 16.1 through the Content Delivery Network (CDN) using **subscription-manager**. For more information, see [Preparing the undercloud](#).

**WARNING**

Some packages in the Red Hat OpenStack Platform software repositories conflict with packages provided by the Extra Packages for Enterprise Linux (EPEL) software repositories. The use of Red Hat OpenStack Platform on systems with the EPEL software repositories enabled is unsupported.

1.10.1. Undercloud repositories

Red Hat OpenStack Platform 16.1 runs on Red Hat Enterprise Linux 8.2. As a result, you must lock the content from these repositories to the respective Red Hat Enterprise Linux version.

**WARNING**

Any repositories outside the ones specified here are not supported. Ensure that you do not enable Extra Packages for Enterprise Linux (EPEL) or the upgrade will fail.

Core repositories

The following table lists core repositories for installing the undercloud.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for x86_64 - BaseOS (RPMs) Extended Update Support (EUS)	rhel-8-for-x86_64-baseos-eus-rpms	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 8 for x86_64 - AppStream (RPMs)	rhel-8-for-x86_64-appstream-eus-rpms	Contains Red Hat OpenStack Platform dependencies.
Red Hat Enterprise Linux 8 for x86_64 - High Availability (RPMs) Extended Update Support (EUS)	rhel-8-for-x86_64-highavailability-eus-rpms	High availability tools for Red Hat Enterprise Linux. Used for Controller node high availability.
Red Hat Ansible Engine 2.9 for RHEL 8 x86_64 (RPMs)	ansible-2.9-for-rhel-8-x86_64-rpms	Ansible Engine for Red Hat Enterprise Linux. Used to provide the latest version of Ansible.
Red Hat Satellite Tools for RHEL 8 Server RPMs x86_64	satellite-tools-6.5-for-rhel-8-x86_64-rpms	Tools for managing hosts with Red Hat Satellite 6.

Name	Repository	Description of requirement
Red Hat OpenStack Platform 16.1 for RHEL 8 (RPMs)	openstack-16.1-for-rhel-8-x86_64-rpms	Core Red Hat OpenStack Platform repository, which contains packages for Red Hat OpenStack Platform director.
Red Hat Fast Datapath for RHEL 8 (RPMS)	fast-datapath-for-rhel-8-x86_64-rpms	Provides Open vSwitch (OVS) packages for OpenStack Platform.

Ceph repositories

The following table lists Ceph Storage related repositories for the undercloud.

Name	Repository	Description of Requirement
Red Hat Ceph Storage Tools 4 for RHEL 8 x86_64 (RPMs)	rhceph-4-tools-for-rhel-8-x86_64-rpms	Provides tools for nodes to communicate with the Ceph Storage cluster. The undercloud requires the ceph-ansible package from this repository if you plan to use Ceph Storage in your overcloud.

IBM POWER repositories

The following table contains a list of repositories for Red Hat Openstack Platform on POWER PC architecture. Use these repositories in place of equivalents in the Core repositories.

Name	Repository	Description of requirement
Red Hat Enterprise Linux for IBM Power, little endian - BaseOS (RPMs)	rhel-8-for-ppc64le-baseos-rpms	Base operating system repository for ppc64le systems.
Red Hat Enterprise Linux 8 for IBM Power, little endian - AppStream (RPMs)	rhel-8-for-ppc64le-appstream-rpms	Contains Red Hat OpenStack Platform dependencies.
Red Hat Enterprise Linux 8 for IBM Power, little endian - High Availability (RPMs)	rhel-8-for-ppc64le-highavailability-rpms	High availability tools for Red Hat Enterprise Linux. Used for Controller node high availability.
Red Hat Ansible Engine 2.8 for RHEL 8 IBM Power, little endian (RPMs)	ansible-2.8-for-rhel-8-ppc64le-rpms	Ansible Engine for Red Hat Enterprise Linux. Provides the latest version of Ansible.

Name	Repository	Description of requirement
Red Hat OpenStack Platform 16.1 for RHEL 8 (RPMs)	openstack-16.1-for-rhel-8-ppc64le-rpms	Core Red Hat OpenStack Platform repository for ppc64le systems.

1.10.2. Overcloud repositories

Red Hat OpenStack Platform 16.1 runs on Red Hat Enterprise Linux 8.2. As a result, you must lock the content from these repositories to the respective Red Hat Enterprise Linux version.



WARNING

Any repositories outside the ones specified here are not supported. Ensure that you do not enable Extra Packages for Enterprise Linux (EPEL) or the upgrade will fail.

Core repositories

The following table lists core repositories for installing the overcloud.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for x86_64 - BaseOS (RPMs) Extended Update Support (EUS)	rhel-8-for-x86_64-baseos-eus-rpms	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 8 for x86_64 - AppStream (RPMs)	rhel-8-for-x86_64-appstream-eus-rpms	Contains Red Hat OpenStack Platform dependencies.
Red Hat Enterprise Linux 8 for x86_64 - High Availability (RPMs) Extended Update Support (EUS)	rhel-8-for-x86_64-highavailability-eus-rpms	High availability tools for Red Hat Enterprise Linux.
Red Hat Ansible Engine 2.9 for RHEL 8 x86_64 (RPMs)	ansible-2.9-for-rhel-8-x86_64-rpms	Ansible Engine for Red Hat Enterprise Linux. Used to provide the latest version of Ansible.
Advanced Virtualization for RHEL 8 x86_64 (RPMs)	advanced-virt-for-rhel-8-x86_64-rpms	Provides virtualization packages for OpenStack Platform.
Red Hat Satellite Tools for RHEL 8 Server RPMs x86_64	satellite-tools-6.5-for-rhel-8-x86_64-rpms	Tools for managing hosts with Red Hat Satellite 6.

Name	Repository	Description of requirement
Red Hat OpenStack Platform 16.1 for RHEL 8 (RPMs)	openstack-16.1-for-rhel-8-x86_64-rpms	Core Red Hat OpenStack Platform repository.
Red Hat Fast Datapath for RHEL 8 (RPMS)	fast-datapath-for-rhel-8-x86_64-rpms	Provides Open vSwitch (OVS) packages for OpenStack Platform.

Ceph repositories

The following table lists Ceph Storage related repositories for the overcloud.

Name	Repository	Description of Requirement
Red Hat Enterprise Linux 8 for x86_64 - BaseOS (RPMs)	rhel-8-for-x86_64-baseos-rpms	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 8 for x86_64 - AppStream (RPMs)	rhel-8-for-x86_64-appstream-rpms	Contains Red Hat OpenStack Platform dependencies.
Red Hat Enterprise Linux 8 for x86_64 - High Availability (RPMs)	rhel-8-for-x86_64-highavailability-rpms	High availability tools for Red Hat Enterprise Linux.
Red Hat Ansible Engine 2.9 for RHEL 8 x86_64 (RPMs)	ansible-2.9-for-rhel-8-x86_64-rpms	Ansible Engine for Red Hat Enterprise Linux. Used to provide the latest version of Ansible.
Red Hat OpenStack Platform 16.1 Director Deployment Tools for RHEL 8 x86_64 (RPMs)	openstack-16.1-deployment-tools-for-rhel-8-x86_64-rpms	Packages to help director configure Ceph Storage nodes.
Red Hat Ceph Storage OSD 4 for RHEL 8 x86_64 (RPMs)	rhceph-4-osd-for-rhel-8-x86_64-rpms	(For Ceph Storage Nodes) Repository for Ceph Storage Object Storage daemon. Installed on Ceph Storage nodes.
Red Hat Ceph Storage MON 4 for RHEL 8 x86_64 (RPMs)	rhceph-4-mon-for-rhel-8-x86_64-rpms	(For Ceph Storage Nodes) Repository for Ceph Storage Monitor daemon. Installed on Controller nodes in OpenStack environments using Ceph Storage nodes.

Name	Repository	Description of Requirement
Red Hat Ceph Storage Tools 4 for RHEL 8 x86_64 (RPMs)	rhceph-4-tools-for-rhel-8-x86_64-rpms	Provides tools for nodes to communicate with the Ceph Storage cluster. This repository should be enabled for all nodes when deploying an overcloud with a Ceph Storage cluster.

Real Time repositories

The following table lists repositories for Real Time Compute (RTC) functionality.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for x86_64 - Real Time (RPMs)	rhel-8-for-x86_64-rt-rpms	Repository for Real Time KVM (RT-KVM). Contains packages to enable the real time kernel. Enable this repository for all Compute nodes targeted for RT-KVM. NOTE: You need a separate subscription to a Red Hat OpenStack Platform for Real Time SKU to access this repository.
Red Hat Enterprise Linux 8 for x86_64 - Real Time for NFV (RPMs)	rhel-8-for-x86_64-nfv-rpms	Repository for Real Time KVM (RT-KVM) for NFV. Contains packages to enable the real time kernel. Enable this repository for all NFV Compute nodes targeted for RT-KVM. NOTE: You need a separate subscription to a Red Hat OpenStack Platform for Real Time SKU to access this repository.

IBM POWER repositories

The following table lists repositories for Openstack Platform on POWER PC architecture. Use these repositories in place of equivalents in the Core repositories.

Name	Repository	Description of requirement
Red Hat Enterprise Linux for IBM Power, little endian - BaseOS (RPMs)	rhel-8-for-ppc64le-baseos-rpms	Base operating system repository for ppc64le systems.
Red Hat Enterprise Linux 8 for IBM Power, little endian - AppStream (RPMs)	rhel-8-for-ppc64le-appstream-rpms	Contains Red Hat OpenStack Platform dependencies.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for IBM Power, little endian - High Availability (RPMs)	rhel-8-for-ppc64le-highavailability-rpms	High availability tools for Red Hat Enterprise Linux. Used for Controller node high availability.
Red Hat Ansible Engine 2.8 for RHEL 8 IBM Power, little endian (RPMs)	ansible-2.8-for-rhel-8-ppc64le-rpms	Ansible Engine for Red Hat Enterprise Linux. Used to provide the latest version of Ansible.
Red Hat OpenStack Platform 16.1 for RHEL 8 (RPMs)	openstack-16.1-for-rhel-8-ppc64le-rpms	Core Red Hat OpenStack Platform repository for ppc64le systems.

1.11. PRODUCT SUPPORT

Available resources include:

Customer Portal

The Red Hat Customer Portal offers a wide range of resources to help guide you through planning, deploying, and maintaining your Red Hat OpenStack Platform deployment. Facilities available via the Customer Portal include:

- Product documentation
- Knowledge base articles and solutions
- Technical briefs
- Support case management

Access the Customer Portal at <https://access.redhat.com/>.

Mailing Lists

Red Hat provides these public mailing lists that are relevant to Red Hat OpenStack Platform users:

- The **rhsa-announce** mailing list provides notification of the release of security fixes for all Red Hat products, including Red Hat OpenStack Platform.

Subscribe at <https://www.redhat.com/mailman/listinfo/rhsa-announce>.

CHAPTER 2. TOP NEW FEATURES

This section provides an overview of the top new features in this release of Red Hat OpenStack Platform.

2.1. COMPUTE

This section outlines the top new features for the Compute service.

Tenant-isolated host aggregates using the Placement service

You can use the Placement service to provide tenant isolation by creating host aggregates that only specific tenants can launch instances on. For more information, see [Creating a tenant-isolated host aggregate](#).

File-backed memory

You can configure instances to use a local storage device as the memory backing device.

2.2. DISTRIBUTED COMPUTE NODES (DCN)

This section outlines the top new features for Distributed Compute Nodes (DCN).

Multi-stack for Distributed Compute Node (DCN)

In Red Hat OpenStack Platform 16.1, you can partition a single overcloud deployment into multiple heat stacks in the undercloud to separate deployment and management operations within a DCN deployment. You can deploy and manage each site in a DCN deployment independently with a distinct heat stack.

2.3. NETWORKING

This section outlines the top new features for the Networking service.

HA support for the Load-balancing service (octavia)

In Red Hat OpenStack Platform 16.1, you can make Load-balancing service (octavia) instances highly available when you implement an active-standby topology and use the amphora provider driver. For more information, see [Enabling Amphora active-standby topology](#) in the *Networking Guide*.

Load-balancing service (octavia) support for UDP traffic

You can use the Red Hat OpenStack Platform Load-balancing service (octavia) to balance network traffic on UDP ports. For more information, see [Creating a UDP load balancer with a health monitor](#) in the *Networking Guide*.

2.4. STORAGE

This section outlines the top new features for the Storage service.

Storage at the Edge with Distributed Compute Nodes (DCN)

In Red Hat OpenStack Platform 16.1, you can deploy storage at the edge with Distributed Compute Nodes. The following features have been added to support this architecture:

- Image Service (glance) multi-stores with RBD.
- Image Service multi-store image import tooling.

- Block Storage Service (cinder) A/A at the edge.
- Support for director deployments with multiple Ceph clusters.

Support for Manila CephFS Native

In Red Hat OpenStack Platform 16.1, the Shared Filesystems service (manila) fully supports the Native CephFS driver.

2.5. BARE METAL SERVICE

This section outlines the top new features for the Bare Metal (ironic) service.

Policy-based routing

With this enhancement, you can use policy-based routing for OpenStack nodes to configure multiple route tables and routing rules with **os-net-config**. Policy-based routing uses route tables where, on a host with multiple links, you can send traffic through a particular interface depending on the source address. You can also define route rules for each interface.

2.6. NETWORK FUNCTIONS VIRTUALIZATION

This section outlines the top new features for Network Functions Virtualization (NFV).

Hyper-converged Infrastructure (HCI) deployments with OVS-DPDK

Red Hat OpenStack Platform 16.1 includes support for hyper-covered infrastructure (HCI) deployments with OVS-DPDK. In an HCI architecture, overcloud nodes with Compute and Ceph Storage services are co-located and configured for optimized resource usage.

Open vSwitch (OVS) hardware offload with OVS-ML2

In Red Hat OpenStack Platform 16.1, the OVS switching function has been offloaded to the SmartNIC hardware. This enhancement reduces the processing resources required, and accelerates the datapath. In Red Hat OpenStack Platform 16.1, this feature has graduated from Technology Preview and is now fully supported.

2.7. TECHNOLOGY PREVIEWS

This section outlines features that are in technology preview in Red Hat OpenStack Platform 16.1 GA.



NOTE

For more information on the support scope for features marked as technology previews, see [Technology Preview Features Support Scope](#).

Persistent memory for instances

As a cloud administrator, you can create and configure persistent memory name spaces on Compute nodes that have NVDIMM hardware. Your cloud users can use these nodes to create instances that use the persistent memory name spaces to provide vPMEM.

Memory encryption for instances

As a cloud administrator, you can now configure SEV-capable Compute nodes to provide cloud users the ability to create instances with memory encryption enabled. For more information, see [Configuring SEV-capable Compute nodes to provide memory encryption for instances](#).

Undercloud minion

This release contains the ability to install undercloud minions. An undercloud minion provides additional **heat-engine** and **ironic-conductor** services on a separate host. These additional services support the undercloud with orchestration and provisioning operations. The distribution of undercloud operations across multiple hosts provides more resources to run an overcloud deployment, which can result in potentially faster and larger deployments.

Deploying bare metal over IPv6 with director

If you have IPv6 nodes and infrastructure, you can configure the undercloud and the provisioning network to use IPv6 instead of IPv4 so that director can provision and deploy Red Hat OpenStack Platform onto IPv6 nodes. For more information, see [Configuring the undercloud for bare metal provisioning over IPv6](#) and [Configuring a custom IPv6 provisioning network](#).

Nova-less provisioning

In Red Hat OpenStack Platform 16.1, you can separate the provisioning and deployment stages of your deployment into distinct steps:

1. Provision your bare metal nodes.
 - a. Create a node definition file in yaml format.
 - b. Run the provisioning command, including the node definition file.
2. Deploy your overcloud.
 - a. Run the deployment command, including the heat environment file that the provisioning command generates.

The provisioning process provisions your nodes and generates a heat environment file that contains various node specifications, including node count, predictive node placement, custom images, and custom NICs. When you deploy your overcloud, include this file in the deployment command.

OVN Load-balancing service (octavia) provider driver

The OVN Load-balancing service provider driver is an integration driver between load balancers provided by OVN and octavia. It supports basic load balancer functionalities and is based on Openflow rules.

The provider driver is automatically enabled in the Load-balancing service by director on OVN Neutron ML2 enabled deployments. There are no additional installation or configuration steps required. The Amphora provider driver remains enabled and is the default provider driver.

Dataplane routed provider networks

Red Hat OpenStack Platform now has official support for routed provider networks with the ML2/OVS mechanism driver. You can use a routed provider network to enable a single provider network to represent multiple layer-2 networks (broadcast domains) or segments so that the operator can present only one network to users. This is a common network type in Edge/DCN deployments. These deployments should include a local DHCP agent, and define a nova availability zone for each edge site or segment.

CHAPTER 3. RELEASE INFORMATION

These release notes highlight technology preview items, recommended practices, known issues, and deprecated functionality that you should consider when you deploy this release of Red Hat OpenStack Platform. Notes for updates released during the support lifecycle of this Red Hat OpenStack Platform release appear in the advisory text associated with each update.

3.1. RED HAT OPENSTACK PLATFORM 16.1 GA

These release notes highlight technology preview items, recommended practices, known issues, and deprecated functionality to be taken into consideration when deploying this release of Red Hat OpenStack Platform.

3.1.1. Bug Fix

These bugs were fixed in this release of Red Hat OpenStack Platform:

BZ#1853275

Before this update, director did not set the **noout** flag on Red Hat Ceph Storage OSDs before running a Leapp upgrade. As a result, additional time was required for the the OSDs to rebalance after the upgrade.

With this update, director sets the **noout** flag before the Leapp upgrade, which accelerates the upgrade process. Director also unsets the **noout** flag after the Leapp upgrade.

BZ#1594033

Before this update, the latest volume attributes were not updated during poll, and the volume data was incorrect on the display screen. With this update, volume attributes update correctly during poll and the correct volume data appears on the display screen.

BZ#1792477

Before this update, the overcloud deployment process did not create the TLS certificate necessary for the Block Storage Service (cinder) to run in active/active mode. As a result, cinder services failed during start-up. With this update, the deployment process creates the TLS certificate correctly and cinder can run in active/active mode with TLS-everywhere.

BZ#1803989

Before this update, it was not possible to deploy the overcloud in a Distributed Compute Node (DCN) or spine-leaf configuration with stateless IPv6 on the control plane. Deployments in this scenario failed during ironic node server provisioning. With this update, you can now deploy successfully with stateless IPv6 on the control plane.

BZ#1804079

Before this update, the etcd service was not configured properly to run in a container. As a result, an error occurred when the service tried to create the TLS certificate. With this update, the etcd service runs in a container and can create the TLS certificate.

BZ#1813391

With this update, PowerMax configuration options are correct for iSCSI and FC drivers. For more information, see <https://docs.openstack.org/cinder/latest/configuration/block-storage/drivers/dell-emc-powermax-driver.html>

BZ#1813393

PowerMax configuration options have changed since Newton. This update includes the latest PowerMax configuration options and supports both iSCSI and FC drivers.

The **CinderPowermaxBackend** parameter also supports multiple back ends.

CinderPowermaxBackendName supports a list of back ends, and you can use the new **CinderPowermaxMultiConfig** parameter to specify parameter values for each back end. For example syntax, see **environments/cinder-dellemc-powermax-config.yaml**.

BZ#1814166

With this update, the Red Hat Ceph Storage dashboard uses Ceph 4.1 and a Grafana container based on **ceph4-rhel8**.

BZ#1815305

Before this update, in DCN + HCI deployments with an IPv6 internal API network, the cinder and etcd services were configured with malformed etcd URIs, and the cinder and etcd services failed on startup.

With this update, the IPv6 addresses in the etcd URI are correct and the cinder and etcd services start successfully.

BZ#1815928

Before this update, in deployments with an IPv6 internal API network, the Block Storage Service (cinder) and Compute Service (nova) were configured with a malformed glance-api endpoint URI. As a result, cinder and nova services located in a DCN or Edge deployment could not access the Image Service (glance).

With this update, the IPv6 addresses in the glance-api endpoint URI are correct and the cinder and nova services at Edge sites can access the Image Service successfully.

BZ#1826741

Before this update, the Block Storage service (cinder) assigned the default volume type in a **volume create** request, ignoring alternative methods of specifying the volume type.

With this update, the Block Storage service performs as expected:

- If you specify a **source_valid** in the request, the volume type that the Block Storage service sets is the volume type of the source volume.
- If you specify a **snapshot_id** in the request, the volume type is inferred from the volume type of the snapshot.
- If you specify an **imageRef** in the request, and the image has a **cinder_img_volume_type** image property, the volume type is inferred from the value of the image property. Otherwise, the the Block Storage service sets the volume type is the default volume type that you configure. If you do not configure a volume type, the Block Storage service uses the system default volume type, **DEFAULT**.

When you specify a volume type explicitly in the **volume create** request, the Block Storage service uses the type that you specify.

BZ#1827721

Before this update, there were no retries and no timeout when downloading a final instance image with the direct deploy interface in ironic. As a result, the deployment could fail if the server that hosts the image fails to respond.

With this update, the image download process attempts 2 retries and has a connection timeout of 60 seconds.

BZ#1831893

A regression was introduced in `ipmitool-1.8.18-11` that caused IPMI access to take over 2 minutes for certain BMCs that did not support the "Get Cipher Suites". As a result, introspection could fail and deployments could take much longer than previously.

With this update, `ipmitool` retries are handled differently, introspection passes, and deployments succeed.



NOTE

This issue with `ipmitool` is resolved in `ipmitool-1.8.18-17`.

BZ#1832720

Before this update, stale **neutron-haproxy-qdhcp-*** containers remained after you deleted the related network. With this update, all related containers are cleaned correctly when you delete a network.

BZ#1832920

Before this update, the **ExtraConfigPre per_node** script was not compatible with Python 3. As a result, the overcloud deployment failed at the step **TASK [Run deployment NodeSpecificDeployment]** with the message **SyntaxError: invalid syntax**.

With this update, the **ExtraConfigPre per_node** script is compatible with Python 3 and you can provision custom **per_node** hieradata.

BZ#1845079

Before this update, the data structure format that the **ceph osd stat -f json** command returns changed. As a result, the validation to stop the deployment unless a certain percentage of Red Hat Ceph Storage (RHCS) OSDs are running did not function correctly, and stopped the deployment regardless of how many OSDs were running.

With this update, the new version of **openstack-tripleo-validations** computes the percentage of running RHCS OSDs correctly and the deployment stops early if a percentage of RHCS OSDs are not running. You can use the parameter **CephOsdPercentageMin** to customize the percentage of RHCS OSDs that must be running. The default value is 66%. Set this parameter to **0** to disable the validation.

BZ#1850991

Before this update, the Red Hat Ceph Storage dashboard listener was created in the HA Proxy configuration, even if the dashboard is disabled. As a result, upgrades of OpenStack with Ceph could fail.

With this update, the service definition has been updated to distinguish the Ceph MGR service from the dashboard service so that the dashboard service is not configured if it is not enabled and upgrades are successful.

BZ#1853433

Before this update, the Leapp upgrade could fail if you had any NFS shares mounted. Specifically, the nodes that run the Compute Service (`nova`) or the Image Service (`glance`) services hung if they used an NFS mount.

With this update, before the Leapp upgrade, director unmounts **/var/lib/nova/instances**, **/var/lib/glance/images**, and any Image Service staging area that you define with the **GlanceNodeStagingUri** parameter.

3.1.2. Enhancements

This release of Red Hat OpenStack Platform features the following enhancements:

BZ#1440926

With this enhancement, you can configure Red Hat OpenStack Platform to use an external, pre-existing Ceph RadosGW cluster. You can manage this cluster externally as an object-store for OpenStack guests.

BZ#1575512

With this enhancement, you can control multicast over the external networks and avoid cluster autoforming over external networks instead of only the internal networks.

BZ#1598716

With this enhancement, you can use director to deploy the Image Service (glance) with multiple image stores. For example, in a Distributed Compute Node (DCN) or Edge deployment, you can store images at each site.

BZ#1617923

With this update, the Validation Framework CLI is fully operational. Specifically, the **openstack tripleo validator** command now includes all of the CLI options necessary to list, run, and show validations, either by validation name or by group.

BZ#1676989

With this enhancement, you can use ATOS HSM deployment with HA mode.

BZ#1686001

With this enhancement, you can revert Block Storage (cinder) volumes to the most recent snapshot, if supported by the driver. This method of reverting a volume is more efficient than cloning from a snapshot and attaching a new volume.

BZ#1698527

With this update, the OVS switching function has been offloaded to the SmartNIC hardware. This enhancement reduces the processing resources required, and accelerates the datapath. In Red Hat OpenStack Platform 16.1, this feature has graduated from Technology Preview and is now fully supported.

BZ#1701416

With this enhancement, HTTP traffic that travels from the HAProxy load balancer to Red Hat Ceph Storage RadosGW instances is encrypted.

BZ#1740946

With this update, you can deploy pre-provisioned nodes with TLSe using the new 'tripleo-ipa' method.

BZ#1767581

With this enhancement, you can use the **--limit**, **--skip-tags**, and **--tags** Ansible options in the **openstack overcloud deploy** command. This is particularly useful when you want to run the deployment on specific nodes, for example, during scale-up operations.

BZ#1793525

When you deploy Red Hat Ceph Storage with director, you can define and configure Ceph device classes and map these classes to specific pools for varying workloads.

BZ#1807841

With this update, the **swift_rsync** container runs in unprivileged mode. This makes the **swift_rsync** container more secure.

BZ#1811490

With this enhancement, there are new options in the **openstack tripleo container image push** command that you can use to provide credentials for the source registry. The new options are **--source-username** and **--source-password**.

Before this update, you could not provide credentials when pushing a container image from a source registry that requires authentication. Instead, the only mechanism to push the container was to pull the image manually and push from the local system.

BZ#1814278

With this enhancement, you can use policy-based routing for Red Hat OpenStack Platform nodes to configure multiple route tables and routing rules with **os-net-config**.

Policy-based routing uses route tables where, on a host with multiple links, you can send traffic through a particular interface depending on the source address. You can also define route rules for each interface.

BZ#1819016

With this update, the **container_images_file** parameter is now a required option in the **undercloud.conf** file. You must set this parameter before you install the undercloud.

With the recent move to use registry.redhat.io as the container source, you must authenticate when you fetch containers. For the undercloud, the **container_images_file** is the recommended option to provide the credentials when you perform the installation. Before this update, if this parameter was not set, the deployment failed with authentication errors when trying to fetch containers.

BZ#1823932

With this enhancement, FreeIPA has DNS entries for the undercloud and overcloud nodes. DNS PTR records are necessary to generate certain types of certificates, particularly certificates for cinder active/active environments with etcd. You can disable this functionality with the **IdMModifyDNS** parameter in an environment file.

BZ#1834185

With this enhancement, you can manage vPMEM with two new parameters **NovaPMEMMappings** and **NovaPMEMNamespaces**.

Use **NovaPMEMMappings** to set the nova configuration option **pmem_namespaces** that reflects mappings between vPMEM and physical PMEM namespaces.

Use **NovaPMEMNamespaces** to create and manage physical PMEM namespaces that you use as a back end for vPMEM.

BZ#1858023

This update includes support for hyper-covered infrastructure (HCI) deployments with OVS-DPDK. In an HCI architecture, overcloud nodes with Compute and Ceph Storage services are co-located and configured for optimized resource usage.

3.1.3. Technology Preview

The items listed in this section are provided as Technology Previews. For further information on the scope of Technology Preview status, and the associated support implications, refer to <https://access.redhat.com/support/offerings/techpreview/>.

BZ#1261083, BZ#1459187

In Red Hat OpenStack Platform 16.1, a technology preview has been added to the Bare Metal Provisioning service (ironic) for deploying the overcloud on an IPv6 provisioning network. For more information, see "Configuring a custom IPv6 provisioning network," in the Bare Metal Provisioning guide.

BZ#1474394

Red Hat OpenStack Platform 16.1 includes support for bare metal provisioning over an IPv6 provisioning network for BMaaS (Bare Metal as-a-Service) tenants.

BZ#1603440

DNS-as-a-Service (designate) returns to technology preview status in Red Hat OpenStack Platform 16.1.

BZ#1623977

In Red Hat OpenStack Platform 16.1, you can configure Load-balancing service (octavia) instances to forward traffic flow and administrative logs from inside the amphora to a syslog server.

BZ#1666684

In Red Hat OpenStack Platform 16.1, a technology preview is available for SR-IOV to work with OVN and the Networking service (neutron) driver without requiring the Networking service DHCP agent. When virtual machines boot on hypervisors that support SR-IOV NICs, the local OVN controllers can reply to the DHCP, internal DNS, and IPv6 router solicitation requests from the virtual machine.

BZ#1671811

In Red Hat OpenStack Platform 16.1 there is a technology preview for routed provider networks with the ML2/OVS mechanism driver. You can use a routed provider network to enable a single provider network to represent multiple layer 2 networks (broadcast domains) or segments so that the operator can present only one network to users. This is a common network type in Edge DCN deployments and Spine-Leaf routed datacenter deployments.

Because Nova scheduler is not segment-aware, you must map each leaf, rack segment, or DCN edge site to a Nova host-aggregate or availability zone. If the deployments require DHCP or the metadata service, you must also define a Nova availability zone for each edge site or segment.

Known Limitations:

- Supported with ML2/OVS only. Not supported with OVN (RFE Bug 1797664)
- Nova scheduler is not segment-aware. For successful nova scheduling, map each segment or edge site to a Nova host-aggregate or availability zone. Currently there are only 2 instance boot options available [RFE Bug 1761903]
 - Boot an instance with port-id and no IP address (differ IP address assignment and specify Nova AZ (segment or edge site))
 - Boot with network-id and specify Nova AZ (segment or edge site)
- Because Nova scheduler is not segment-aware, Cold/Live migration works only when you specify the destination Nova availability zone (segment or edge site) [RFE Bug 1761903]
- North-south routing with central SNAT or Floating IP is not supported [RFE Bug 1848474]
- When using SR-IOV or PCI pass-through, physical network (physnet) names must be the same in central and remote sites or segments. You cannot reuse segment-ids (Bug 1839097)

For more information, see <https://docs.openstack.org/neutron/train/admin/config-routed-networks.html>.

BZ#1676631

In Red Hat OpenStack Platform 16.1, the Open Virtual Network (OVN) provider driver for the Load-balancing service (octavia) is in technology preview.

BZ#1703958

This update includes support for both TCP and UDP protocols on the same load-balancer listener for OVN Provider driver.

BZ#1758424

With this update, when using Image Service (glance) multi stores, the image owner can delete an Image copy from a specific store.

BZ#1801721

In Red Hat OpenStack Platform 16.1, the Load-balancing service (Octavia) has a technology preview for UDP protocol.

BZ#1848582

With this release, a technology preview has been added for the Shared File Systems service (manila) for IPv6 to work in the CephFS NFS driver. This feature requires Red Hat Ceph Storage 4.1.

3.1.4. Rebase: Bug Fixes and Enhancements

These items are rebases of bug fixes and enhancements included in this release of Red Hat OpenStack Platform:

BZ#1738449

collectd 5.11 contains bug fixes and new plugins. For more information, see <https://github.com/collectd/collectd/releases>.

3.1.5. Release Notes

This section outlines important details about the release, including recommended practices and notable changes to Red Hat OpenStack Platform. You must take this information into account to ensure the best possible outcomes for your deployment.

BZ#1225775

The Image Service (glance) now supports multi stores with the Ceph RBD driver.

BZ#1546996

With this release, **networking-ovn** now supports QoS bandwidth limitation and DSCP marking rules with the neutron QoS API.

BZ#1654408

For glance image conversion, the **glance-direct** method is not enabled by default. To enable this feature, set **enabled_import_methods** to **[glance-direct,web-download]** or **[glance-direct]** in the **DEFAULT** section of the **glance-api.conf** file.

The Image Service (glance) must have a staging area when you use the **glance-direct** import method. Set the **node_staging_uri** option in the **DEFAULT** section of the **glance-api.conf** file to [file://<absolute-directory-path>](#). This path must be on a shared file system that is available to all Image Service API nodes.

BZ#1700402

Director can now deploy the Block Storage Service in an active/active mode. This deployment scenario is supported only for Edge use cases.

BZ#1710465

When you upgrade from Red Hat OpenStack Platform (RHOSP) 13 DCN to RHOSP 16.1 DCN, it is not possible to migrate from the single stack RHOSP 13 deployment into a multi-stack RHOSP 16.1 deployment. The RHOSP 13 stack continues to be managed as a single stack in the Orchestration service (heat) even after you upgrade to RHOSP 16.1.

After you upgrade to RHOSP 16.1, you can deploy new DCN sites as new stacks. For more information, see the multi-stack documentation for RHOSP 16.1 DCN.

BZ#1758416

In Red Hat OpenStack Platform 16.1, you can use the Image service (glance) to copy existing image data into multiple stores with a single command. This removes the need for the operator to copy data manually and update image locations.

BZ#1758420

In Red Hat OpenStack Platform 16.1, you can use the Image Service (glance) to copy existing image data into multiple stores with a single command. This removes the need for the operator to copy data manually and update image locations.

BZ#1784640

Before this update, during Red Hat Ceph Storage (RHCS) deployment, Red Hat OpenStack Platform (RHOSP) director generated the CephClusterFSID by passing the desired FSID to `ceph-analysable` and used the Python `uuid1()` function. With this update, director uses the Python `uuid4()` function, which generates UUIDs more randomly.

BZ#1790756

With this release, a new feature has been added for the Shared File Systems service (manila) for IPv6 to work in the CephFS NFS driver. This feature requires Red Hat Ceph Storage 4.1.

BZ#1808583

Red Hat OpenStack Platform 16.1 includes the following PowerMax Driver updates:
Feature updates:

- PowerMax Driver - Unisphere storage group/array tagging support
- PowerMax Driver - Short host name and port group name override
- PowerMax Driver - SRDF Enhancement
- PowerMax Driver - Support of Multiple Replication
- Bug fixes:
 - PowerMax Driver - Debug Metadata Fix
 - PowerMax Driver - Volume group delete failure
 - PowerMax Driver - Setting minimum Unisphere version to 9.1.0.5
 - PowerMax Driver - Unmanage Snapshot Delete Fix
 - PowerMax Driver - RDF clean snapvx target fix
 - PowerMax Driver - Get Manageable Volumes Fix
 - PowerMax Driver - Print extend volume info
 - PowerMax Driver - Legacy volume not found
 - PowerMax Driver - Safeguarding retype to some in-use replicated modes
 - PowerMax Driver - Replication array serial check
 - PowerMax Driver - Support of Multiple Replication

- PowerMax Driver - Update single underscores
- PowerMax Driver - SRDF Replication Fixes
- PowerMax Driver - Replication Metadata Fix
- PowerMax Driver - Limit replication devices
- PowerMax Driver - Allowing for default volume type in group
- PowerMax Driver - Version comparison correction
- PowerMax Driver - Detach RepConfig logging & Retype rename remote fix
- PowerMax Driver - Manage volume emulation check
- PowerMax Driver - Deletion of group with volumes
- PowerMax Driver - PowerMax Pools Fix
- PowerMax Driver - RDF status validation
- PowerMax Driver - Concurrent live migrations failure
- PowerMax Driver - Live migrate remove rep vol from sg
- PowerMax Driver - U4P failover lock not released on exception
- PowerMax Driver - Compression Change Bug Fix

BZ#1810045

The Shared Filesystems service (manila) fully supports the Native CephFS driver. This driver was previously in Tech Preview status, but is now fully supported.

BZ#1846039

The **sg-bridge** container uses the **sg-bridge** RPM to provide an AMQP1-to-unix socket interface for sg-core. Both components are part of the Service Telemetry Framework. This is the initial release of the **sg-bridge** component.

BZ#1852084

Red Hat OpenStack Platform 16.1 includes tripleo-heat-templates support for VXFlexOS Volume Backend.

BZ#1852087

Red Hat OpenStack Platform 16.1 includes support for SC Cinder Backend. The SC Cinder back end now supports both iSCSI and FC drivers, and can also support multiple back ends. You can use the **CinderScBackendName** parameter to list back ends, and the **CinderScMultiConfig** parameter to specify parameter values for each back end. For an example configuration file, see **environments/cinder-dellemc-sc-config.yaml**.

BZ#1855096

The NetApp Backend Guide for the Shared File Systems service (manila) has been removed from the Red Hat OpenStack product documentation pages. This content is now hosted within the NetApp OpenStack documentation suite: https://netapp-openstack-dev.github.io/openstack-docs/train/manila/configuration/manila_config_files/section_rhosp_director_configuration.html

BZ#1858352

If you want to upgrade from Red Hat OpenStack Platform (RHOSP) 13 and Red Hat Ceph Storage (RHCS) 3 with filestore to RHOSP 16.1 and RHCS 4, you cannot migrate to bluestore after the upgrade. You can run RHCS 4 with filestore until a fix is available. For more information, see https://bugzilla.redhat.com/show_bug.cgi?id=1854973.

BZ#1858938

The **sg-bridge** and **sg-core** container images provide a new data path for collectd metrics into the Service Telemetry Framework.

The **sg-bridge** component provides an AMQP1 to unix socket translation to the **sg-core**, resulting in a 500% performance increase over the legacy Smart Gateway component.

This is the initial release of the sg-bridge and sg-core container image components.



NOTE

The legacy Smart Gateway is still the data path for Ceilometer metrics, Ceilometer events, and collectd events.

3.1.6. Known Issues

These known issues exist in Red Hat OpenStack Platform at this time:

BZ#1508449

OVN serves DHCP as an openflow controller with ovn-controller directly on Compute nodes. However, SR-IOV instances are attached directly to the network through the VF/PF and so SR-IOV instances cannot receive DHCP responses.

Workaround: Change **OS::TripleO::Services::NeutronDhcpAgent** to **OS::TripleO::Services::NeutronDhcpAgent: deployment/neutron/neutron-dhcp-container-puppet.yaml**.

BZ#1574431

Currently, quota commands do not work as expected in the Block Storage service (cinder). With the Block Storage CLI, you can successfully create quota entries and the CLI does not check for a valid project ID. Quota entries that the CLI creates without valid project IDs are dummy records that contain invalid data. Until this issue is fixed, if you are a CLI user, you must specify a valid project ID when you create quota entries and monitor Block Storage for dummy records.

BZ#1797047

The Shared File System service (manila) access-list feature requires Red Hat Ceph Storage (RHCS) 4.1 or later. RHCS 4.0 has a packaging issue that means you cannot use the Shared File System service access-list with RHCS 4.0. You can still use share creation, however, the share is unusable without access-list. Consequently, customers who use RHCS 4.0 cannot use the Shared File System service with CephFS via NFS. For more information, see https://bugzilla.redhat.com/show_bug.cgi?id=1797075.

BZ#1828889

There is a known issue where the OVN mechanism driver does not use the Networking Service (neutron) database, but relies on the OVN database instead. As a result, the SR-IOV agent is registered in the Networking Service database because it is outside of OVN. There is currently no workaround for this issue.

BZ#1836963

Because of a core OVN bug, virtual machines with floating IP (FIP) addresses cannot route to other networks in an ML2/OVN deployment with distributed virtual routing (DVR) enabled. Core OVN sets

a bad next hop when routing SNAT IPv4 traffic from a VM with a floating ip with DVR enabled. Instead of the gateway IP, OVN sets the destination IP. As a result, the router sends an ARP request for an unknown IP instead of routing the request to the gateway.

Workaround: Before you deploy a new overcloud with ML2/OVN, disable DVR by setting **NeutronEnableDVR: false** in an environment file. If you have ML2/OVN in an existing deployment, complete the following steps:

1) Set `enable_distributed_floating_ips` to 'False' in the **neutron.conf** file:

```
(undercloud) [stack@undercloud-0 ~]$ ansible -i /usr/bin/tripleo-ansible-inventory -m shell -b -a
"crudini --set /var/lib/config-data/puppet-generated/neutron/etc/neutron/plugins/ml2/ml2_conf.ini ovn enable_distributed_floating_ip False"
Controller
```

2) Restart neutron server containers:

```
(undercloud) [stack@undercloud-0 ~]$ ansible -i /usr/bin/tripleo-ansible-inventory -m shell -b -a
"podman restart neutron_api" Controller
```

3) Centralize all of the FIP traffic through gateway nodes. Run the following command on any overcloud node:

```
$ export NB=$(sudo ovs-vsctl get open . external_ids:ovn-remote | sed -e 's/\\"//g' | sed -e
's/6642/6641/g') $ alias ovn-nbctl='sudo podman exec ovn_controller ovn-nbctl --db=$NB' $ for fip
in $(ovn-nbctl --bare --columns _uuid find nat type=dnat_and_snat); do ovn-nbctl clear NAT $fip
external_mac; done
```

When the fix is available in RHOSP 16.1.1, you can re-enable distributed FIP traffic:

1) Set **enable_distributed_floating_ips** back to 'True' in the **neutron.conf** file:


```
(undercloud) [stack@undercloud-0 ~]$ ansible -i /usr/bin/tripleo-ansible-inventory -m shell -b -a
"crudini --set /var/lib/config-data/puppet-generated/neutron/etc/neutron/plugins/ml2/ml2_conf.ini ovn enable_distributed_floating_ip True"
Controller
```

2) Restart neutron server containers:

```
(undercloud) [stack@undercloud-0 ~]$ ansible -i /usr/bin/tripleo-ansible-inventory -m shell -b -a
"podman restart neutron_api" Controller
```

3) Trigger the update in all of the FIPs. Run the following command on any overcloud node:

```
$ export NB=$(sudo ovs-vsctl get open . external_ids:ovn-remote | sed -e 's/\\"//g' | sed -e
's/6642/6641/g') $ alias ovn-nbctl='sudo podman exec ovn_controller ovn-nbctl --db=$NB' $ for i in
$(ovn-nbctl --bare --columns logical_port find nat type=dnat_and_snat); do ovn-nbctl set
logical_switch_port $i up=false; done
```



NOTE

Disabling DVR causes traffic to be centralized. All L3 traffic travels through the Controller/Networker nodes. This might affect scale, data plane performance, and throughput.

The keepalived instance in the Red Hat OpenStack Platform Load-balancing service (octavia) instance (amphora) can abnormally terminate and interrupt UDP traffic. The cause of this issue is that the timeout value for the UDP health monitor is too small.

Workaround: specify a new timeout value that is greater than two seconds: **\$ openstack loadbalancer healthmonitor set --timeout 3 <health_monitor_id>**

For more information, search for "loadbalancer healthmonitor" in the Command Line Interface Reference.

BZ#1840640

There is an incomplete definition for TLS in the Orchestration service (heat) when you update from 16.0 to 16.1, and the update fails.

To prevent this failure, you must set the following parameter and value: **InternalTLSCAFile: "**.

BZ#1845091

There is a known issue when you update from 16.0 to 16.1 with Public TLS or TLS-Everywhere.

The parameter **InternalTLSCAFile** provides the location of the CA cert bundle for the overcloud instance. Upgrades and updates fail if this parameter is not set correctly. With new deployments, heat sets this parameter correctly, but if you upgrade a deployment that uses old heat templates, then the defaults might not be correct.

Workaround: Set the **InternalTLSCAFile** parameter to an empty string "" so that the undercloud uses the certificates in the default trust store.

BZ#1846557


There is a known issue when upgrading from RHOSP 13 to RHOSP 16.1. The value of **HostnameFormatDefault** has changed from **%stackname%-compute-%index%** to **%stackname%-novacompute-%index%**. This change in default value can result in duplicate service entries and have further impacts on operations such as live migration.

Workaround: If you upgrade from RHOSP 13 to RHOSP 16.1, you must override the **HostnameFormatDefault** value to configure the previous default value to ensure that the previous hostname format is retained. If you upgrade from RHOSP 15 or RHOSP 16.0, no action is required.

BZ#1847463

The output format of **tripleo-ansible-inventory** changed in RHOSP 16.1. As a result, the **generate-inventory** step fails.

Workaround: Create the inventory manually.



NOTE

It is not possible to migrate from ML2/OVS to ML2/OVN in RHOSP 16.1.

BZ#1848180

There is a known issue where a heat parameter **InternalTLSCAFile** is used during deployment when the undercloud contacts the external (public) endpoint to create initial resources and projects. If the internal and public interfaces have certificates from different Certificate Authorities (CAs), the deployment fails. Either the undercloud fails to contact the keystone public interface, or the internal interfaces receive malformed configuration.

This scenario affects deployments with TLS Everywhere, when the IPA server supplies the internal interfaces but the public interfaces have a certificate that the operator supplies. This also prevents 'brown field' deployments, where deployments with existing public certificates attempt to redeploy and configure TLS Everywhere.

There is currently no workaround for this defect.

BZ#1848462

Currently, on ML2/OVS and DVR configurations, Open vSwitch (OVS) routes ICMPv6 traffic incorrectly, causing network outages on tenant networks. At this time, there is no workaround for this issue. If you have clouds that rely heavily on IPv6 and might experience issues caused by blocked ICMP traffic, such as pings, do not update to RHOSP 16.1 until this issue is fixed.

BZ#1849235

If you do not set the **UpgradeLevelNovaCompute** parameter to "", live migrations are not possible when you upgrade from RHOSP 13 to RHOSP 16.

BZ#1850192

There is a known issue in the Block Storage Service (cinder) due to the following conditions:

- Red Hat OpenStack Platform 16.1 supports running the cinder-volume service in active/active (A/A) mode at DCN/Edge sites. The control plane still runs active/passive under pacemaker.
- When running A/A, cinder uses the tripleo etcd service for its lock manager.
- When the deployment includes TLS-everywhere (TLS-e), internal API traffic between cinder and etcd, as well as the etcd inter-node traffic should use TLS. RHOSP 16.1 does not support TLS-e in a way that supports the Block Storage Service and etcd with TLS. However, you can configure etcd not to use TLS, even if you configure and enable TLS-e. As a result, TLS is everywhere except for etcd traffic:
- TLS-Everywhere protects traffic in the Block Storage Service
- Only the traffic between the Block Storage Service and etcd, and the etcd inter-node traffic is not protected
- The traffic is limited to Block Storage Service use of etcd for its Distributed Lock Manager (DLM). This traffic contains reference to Block Storage Service object IDs, for example, volume IDs and snapshot IDs, but does not contain any user or tenant credentials. This limitation will be removed in a RHOSP 16.1 update. For more information, see BZ#1848153.

BZ#1852541

There is a known issue with the Object Storage service (swift). If you use pre-deployed nodes, you might encounter the following error message in /var/log/containers/stdouts/swift_rsync.log: "failed to create pid file /var/run/rsyncd.pid: File exists"

Workaround: Enter the following command on all Controller nodes that are pre-deployed:

```
for d in $(podman inspect swift_rsync | jq '.[].GraphDriver.Data.UpperDir') /var/lib/config-  
data/puppet-generated/swift; do sed -i -e '/pid file/d' $d/etc/rsyncd.conf; done
```

BZ#1852801

When you update or upgrade **python3-tripleoclient**, Ansible does not receive the update or upgrade and Ansible or **ceph-ansible** tasks fail.

When you update or upgrade, ensure that Ansible also receives the update so that playbook tasks can run successfully.

BZ#1854334

There is a known issue with the OVN filter packets that **ovn-controller** generates. Router Advertisements that receive ACL processing in OVN are dropped if there is no explicit ACL rule to allow this traffic.

Workaround: Enter the following command to create a security rule:

```
openstack security group rule create --ethertype IPv6 --protocol icmp --icmp-type 134  
<SECURITY_GROUP>
```

BZ#1855423, BZ#1856901

There are some known limitations for Mellanox ConnectX-5 adapter cards in VF LAG mode in OVS OFFLOAD deployments, SRIOV Switchdev mode.

You might encounter the following known issues and limitations when you use the Mellanox ConnectX-5 adapter cards with the virtual function (VF) link aggregation group (LAG) configuration in an OVS OFFLOAD deployment, SRIOV Switchdev mode:

- When at least one VF of any physical function (PF) is still bound or attached to a virtual machine (VM), an internal firmware error occurs when attempting to disable single-root input/output virtualization (SR-IOV) and when unbinding PF using a function such as **ifdown** and **ip link**. To work around the problem, unbind or detach VFs before you perform these actions:
 1. Shut down and detach any VMs.
 2. Remove VF LAG BOND interface from OVS.
 3. Unbind each configured VF: **# echo <VF PCIe BDF> > /sys/bus/pci/drivers/mlx5_core/unbind**
 4. Disable SR-IOV for each PF: **# echo 0 > /sys/class/net/<PF>/device/sriov_numvfs**
- When the **NUM_OF_VFS** parameter configured in the Firmware configuration (using the **mstconfig** tool) is higher than 64, VF LAG mode while deploying OVS OFFLOAD, SRIOV switchdev mode is not supported. Currently, there is no workaround available.

BZ#1856999

The Ceph Dashboard currently does not work with the TLS Everywhere framework because the **dashboard_protocol** parameter was incorrectly omitted from the heat template. As a result, back ends fail to appear when HAproxy is started.

As a temporary solution, create a new environment file that contains the **dashboard_protocol** parameter and include the environment file in your overcloud deployment with the **-e** option:

```
parameter_defaults:
  CephAnsibleExtraConfig:
    dashboard_protocol: 'https'
```

This solution introduces ceph-ansible bug. For more information, see https://bugzilla.redhat.com/show_bug.cgi?id=1860815.

BZ#1859702

There is a known issue where, after an ungraceful shutdown, Ceph containers might not start automatically on system reboot.

Workaround: Remove the old container IDs manually with the **podman rm** command. For more information, see https://bugzilla.redhat.com/show_bug.cgi?id=1858865#c2.

BZ#1861363

OSP 16.0 introduced full support for live migration of pinned instances. Due to a bug in this feature, instances with a real-time CPU policy and more than one real-time CPU cannot migrate successfully. As a result, live migration of real-time instances is not possible. There is currently no workaround.

BZ#1861370

There is a known issue where enabling the **realtime-virtual-host** tuned profile inside guest virtual machines degrades throughput and displays non-deterministic performance. **ovs-dpdk** PMDs are pinned incorrectly to housekeeping CPUs.

Workaround: Use the **cpu-partitioning** tuned profile inside guest virtual machines, write a post-deployment script to update the **tuned.conf** file, and reboot the node:

```
ps_blacklist=ksoftirqd.*;rcuc.*;rcub.*;ktimersoftd.*;.*pmd.*;.*PMD.*;.*DPDK;.*qemu-kvm.*
```

3.1.7. Removed Functionality

BZ#1832405

In this release of Red Hat OpenStack Platform, you can no longer customize the Red Hat Ceph Storage cluster admin keyring secret. Instead, the admin keyring secret is generated randomly during initial deployment.

CHAPTER 4. TECHNICAL NOTES

This chapter supplements the information contained in the text of Red Hat OpenStack Platform "Train" errata advisories released through the Content Delivery Network.

4.1. RHEA-2020:3148 RED HAT OPENSTACK PLATFORM 16.1 GENERAL AVAILABILITY ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2020:3148. Further information about this advisory is available at link: <https://access.redhat.com/errata/RHBA-2020:3148.html>.

Changes to the ansible-role-atos-hsm component:

- With this enhancement, you can use ATOS HSM deployment with HA mode. (BZ#1676989)

Changes to the collectd component:

- collectd 5.11 contains bug fixes and new plugins. For more information, see <https://github.com/collectd/collectd/releases>. (BZ#1738449)

Changes to the openstack-cinder component:

- With this enhancement, you can revert Block Storage (cinder) volumes to the most recent snapshot, if supported by the driver. This method of reverting a volume is more efficient than cloning from a snapshot and attaching a new volume. (BZ#1686001)
- Director can now deploy the Block Storage Service in an active/active mode. This deployment scenario is supported only for Edge use cases. (BZ#1700402)
- This update includes the following enhancements:
 - Support for revert-to-snapshot in VxFlex OS driver
 - Support for volume migration in VxFlex OS driver
 - Support for OpenStack volume replication v2.1 in VxFlex OS driver
 - Support for VxFlex OS 3.5 in the VxFlex OS driver

Changes to the openstack-designate component:

- DNS-as-a-Service (designate) returns to technology preview status in Red Hat OpenStack Platform 16.1. (BZ#1603440)

Changes to the openstack-glance component:

- The Image Service (glance) now supports multi stores with the Ceph RBD driver. (BZ#1225775)
- In Red Hat OpenStack Platform 16.1, you can use the Image service (glance) to copy existing image data into multiple stores with a single command. This removes the need for the operator to copy data manually and update image locations. (BZ#1758416)
- In Red Hat OpenStack Platform 16.1, you can use the Image Service (glance) to copy existing image data into multiple stores with a single command. This removes the need for the operator to copy data manually and update image locations. (BZ#1758420)

- With this update, when using Image Service (glance) multi stores, the image owner can delete an Image copy from a specific store. (BZ#1758424)

Changes to the openstack-ironic component:

- A regression was introduced in ipmitool-1.8.18-11 that caused IPMI access to take over 2 minutes for certain BMCs that did not support the "Get Cipher Suites". As a result, introspection could fail and deployments could take much longer than previously.
With this update, ipmitool retries are handled differently, introspection passes, and deployments succeed.



NOTE

This issue with ipmitool is resolved in ipmitool-1.8.18-17. (BZ#1831893)

Changes to the openstack-ironic-python-agent component:

- Before this update, there were no retries and no timeout when downloading a final instance image with the direct deploy interface in ironic. As a result, the deployment could fail if the server that hosts the image fails to respond.
With this update, the image download process attempts 2 retries and has a connection timeout of 60 seconds. (BZ#1827721)

Changes to the openstack-neutron component:

- Before this update, it was not possible to deploy the overcloud in a Distributed Compute Node (DCN) or spine-leaf configuration with stateless IPv6 on the control plane. Deployments in this scenario failed during ironic node server provisioning. With this update, you can now deploy successfully with stateless IPv6 on the control plane. (BZ#1803989)

Changes to the openstack-tripleo-common component:

- When you update or upgrade **python3-tripleoclient**, Ansible does not receive the update or upgrade and Ansible or **ceph-ansible** tasks fail.
When you update or upgrade, ensure that Ansible also receives the update so that playbook tasks can run successfully. (BZ#1852801)
- With this update, the Red Hat Ceph Storage dashboard uses Ceph 4.1 and a Grafana container based on **ceph4-rhel8**. (BZ#1814166)
- Before this update, during Red Hat Ceph Storage (RHCS) deployment, Red Hat OpenStack Platform (RHOSP) director generated the CephClusterFSID by passing the desired FSID to ceph-ansible and used the Python uuid1() function. With this update, director uses the Python uuid4() function, which generates UUIDs more randomly. (BZ#1784640)

Changes to the openstack-tripleo-heat-templates component:

- There is an incomplete definition for TLS in the Orchestration service (heat) when you update from 16.0 to 16.1, and the update fails.
To prevent this failure, you must set the following parameter and value: **InternalTLSCAFile: "**. (BZ#1840640)
- With this enhancement, you can configure Red Hat OpenStack Platform to use an external, pre-existing Ceph RadosGW cluster. You can manage this cluster externally as an object-store for OpenStack guests. (BZ#1440926)

- With this enhancement, you can use director to deploy the Image Service (glance) with multiple image stores. For example, in a Distributed Compute Node (DCN) or Edge deployment, you can store images at each site. (BZ#1598716)
- With this enhancement, HTTP traffic that travels from the HAProxy load balancer to Red Hat Ceph Storage RadosGW instances is encrypted. (BZ#1701416)
- With this update, you can deploy pre-provisioned nodes with TLSe using the new 'tripleo-ipa' method. (BZ#1740946)
- Before this update, in deployments with an IPv6 internal API network, the Block Storage Service (cinder) and Compute Service (nova) were configured with a malformed glance-api endpoint URI. As a result, cinder and nova services located in a DCN or Edge deployment could not access the Image Service (glance).
With this update, the IPv6 addresses in the glance-api endpoint URI are correct and the cinder and nova services at Edge sites can access the Image Service successfully. (BZ#1815928)
- With this enhancement, FreeIPA has DNS entries for the undercloud and overcloud nodes. DNS PTR records are necessary to generate certain types of certificates, particularly certificates for cinder active/active environments with etcd. You can disable this functionality with the **IdMModifyDNS** parameter in an environment file. (BZ#1823932)
- In this release of Red Hat OpenStack Platform, you can no longer customize the Red Hat Ceph Storage cluster admin keyring secret. Instead, the admin keyring secret is generated randomly during initial deployment. (BZ#1832405)
- Before this update, stale **neutron-haproxy-qdhcp-*** containers remained after you deleted the related network. With this update, all related containers are cleaned correctly when you delete a network. (BZ#1832720)
- Before this update, the **ExtraConfigPre per_node** script was not compatible with Python 3. As a result, the overcloud deployment failed at the step **TASK [Run deployment NodeSpecificDeployment]** with the message **SyntaxError: invalid syntax**.
With this update, the **ExtraConfigPre per_node** script is compatible with Python 3 and you can provision custom **per_node** hieradata. (BZ#1832920)
- With this update, the **swift_rsyc** container runs in unprivileged mode. This makes the **swift_rsyc** container more secure. (BZ#1807841)
- PowerMax configuration options have changed since Newton. This update includes the latest PowerMax configuration options and supports both iSCSI and FC drivers.
The **CinderPowermaxBackend** parameter also supports multiple back ends.
CinderPowermaxBackendName supports a list of back ends, and you can use the new **CinderPowermaxMultiConfig** parameter to specify parameter values for each back end. For example syntax, see **environments/cinder-dellemc-powermax-config.yaml**. (BZ#1813393)
- Support for Xtremio Cinder Backend
Updated the Xtremio cinder backend to support both iSCSI and FC drivers. It is also enhanced to support multiple backends. (BZ#1852082)
- Red Hat OpenStack Platform 16.1 includes tripleo-heat-templates support for VXFlexOS Volume Backend. (BZ#1852084)
- Red Hat OpenStack Platform 16.1 includes support for SC Cinder Backend. The SC Cinder back end now supports both iSCSI and FC drivers, and can also support multiple back ends. You can use the **CinderScBackendName** parameter to list back ends, and the **CinderScMultiConfig**

parameter to specify parameter values for each back end. For an example configuration file, see **environments/cinder-dellemc-sc-config.yaml**. (BZ#1852087)

- PowerMax configuration options have changed since Newton. This update includes the latest PowerMax configuration options and supports both iSCSI and FC drivers. The **CinderPowermaxBackend** parameter also supports multiple back ends. **CinderPowermaxBackendName** supports a list of back ends, and you can use the new **CinderPowermaxMultiConfig** parameter to specify parameter values for each back end. For example syntax, see **environments/cinder-dellemc-powermax-config.yaml**. (BZ#1852088)

Changes to the openstack-tripleo-validations component:

- Before this update, the data structure format that the **ceph osd stat -f json** command returns changed. As a result, the validation to stop the deployment unless a certain percentage of Red Hat Ceph Storage (RHCS) OSDs are running did not function correctly, and stopped the deployment regardless of how many OSDs were running. With this update, the new version of **openstack-tripleo-validations** computes the percentage of running RHCS OSDs correctly and the deployment stops early if a percentage of RHCS OSDs are not running. You can use the parameter **CephOsdPercentageMin** to customize the percentage of RHCS OSDs that must be running. The default value is 66%. Set this parameter to **0** to disable the validation. (BZ#1845079)

Changes to the puppet-cinder component:

- With this update, PowerMax configuration options are correct for iSCSI and FC drivers. For more information, see <https://docs.openstack.org/cinder/latest/configuration/block-storage/drivers/dell-emc-powermax-driver.html> (BZ#1813391)

Changes to the puppet-tripleo component:

- Before this update, the etcd service was not configured properly to run in a container. As a result, an error occurred when the service tried to create the TLS certificate. With this update, the etcd service runs in a container and can create the TLS certificate. (BZ#1804079)

Changes to the python-cinderclient component:

- Before this update, the latest volume attributes were not updated during poll, and the volume data was incorrect on the display screen. With this update, volume attributes update correctly during poll and the correct volume data appears on the display screen. (BZ#1594033)

Changes to the python-networking-ovn component:

- Because of a core OVN bug, virtual machines with floating IP (FIP) addresses cannot route to other networks in an ML2/OVN deployment with distributed virtual routing (DVR) enabled. Core OVN sets a bad next hop when routing SNAT IPv4 traffic from a VM with a floating ip with DVR enabled. Instead of the gateway IP, OVN sets the destination IP. As a result, the router sends an ARP request for an unknown IP instead of routing the request to the gateway. Workaround: Before you deploy a new overcloud with ML2/OVN, disable DVR by setting **NeutronEnableDVR: false** in an environment file. If you have ML2/OVN in an existing deployment, complete the following steps:

1) Set `enable_distributed_floating_ips` to 'False' in the **neutron.conf** file:

```
(undercloud) [stack@undercloud-0 ~]$ ansible -i /usr/bin/tripleo-ansible-inventory -m shell -b -a "crudini --set /var/lib/config-data/puppet-generated/neutron/etc/neutron/plugins/ml2/ml2_conf.ini ovn enable_distributed_floating_ip False" Controller
```

2) Restart neutron server containers:

```
(undercloud) [stack@undercloud-0 ~]$ ansible -i /usr/bin/tripleo-ansible-inventory -m shell -b
-a "podman restart neutron_api" Controller
```

3) Centralize all of the FIP traffic through gateway nodes. Run the following command on any overcloud node:

```
$ export NB=$(sudo ovs-vsctl get open . external_ids:ovn-remote | sed -e 's/\\"//g' | sed -e
's/6642/6641/g') $ alias ovn-nbctl='sudo podman exec ovn_controller ovn-nbctl --db=$NB' $
for fip in $(ovn-nbctl --bare --columns _uuid find nat type=dnat_and_snat); do ovn-nbctl clear
NAT $fip external_mac; done
```

When the fix is available in RHOSP 16.1.1, you can re-enable distributed FIP traffic:

1) Set **enable_distributed_floating_ips** back to 'True' in the **neutron.conf** file:

```
(undercloud) [stack@undercloud-0 ~]$ ansible -i /usr/bin/tripleo-ansible-inventory -m shell -b
-a "crudini --set /var/lib/config-data/puppet-generated/neutron/etc/neutron/plugins/ml2/ml2_conf.ini ovn enable_distributed_floating_ip
True" Controller
```

2) Restart neutron server containers:

```
(undercloud) [stack@undercloud-0 ~]$ ansible -i /usr/bin/tripleo-ansible-inventory -m shell -b
-a "podman restart neutron_api" Controller
```

3) Trigger the update in all of the FIPs. Run the following command on any overcloud node:

```
$ export NB=$(sudo ovs-vsctl get open . external_ids:ovn-remote | sed -e 's/\\"//g' | sed -e
's/6642/6641/g') $ alias ovn-nbctl='sudo podman exec ovn_controller ovn-nbctl --db=$NB' $
for i in $(ovn-nbctl --bare --columns logical_port find nat type=dnat_and_snat); do ovn-nbctl
set logical_switch_port $i up=false; done
```



NOTE

Disabling DVR causes traffic to be centralized. All L3 traffic travels through the Controller/Networker nodes. This might affect scale, data plane performance, and throughput. (BZ#1836963)

Changes to the python-tripleoclient component:

- With this enhancement, you can use the **--limit**, **--skip-tags**, and **--tags** Ansible options in the **openstack overcloud deploy** command. This is particularly useful when you want to run the deployment on specific nodes, for example, during scale-up operations. (BZ#1767581)
- With this enhancement, there are new options in the **openstack tripleo container image push** command that you can use to provide credentials for the source registry. The new options are **--source-username** and **--source-password**.
Before this update, you could not provide credentials when pushing a container image from a source registry that requires authentication. Instead, the only mechanism to push the container was to pull the image manually and push from the local system. (BZ#1811490)
- With this update, the **container_images_file** parameter is now a required option in the **undercloud.conf** file. You must set this parameter before you install the undercloud. With the recent move to use registry.redhat.io as the container source, you must authenticate

when you fetch containers. For the undercloud, the **container_images_file** is the recommended option to provide the credentials when you perform the installation. Before this update, if this parameter was not set, the deployment failed with authentication errors when trying to fetch containers. (BZ#1819016)