Red Hat OpenStack Platform 16.1

Release Notes

Release details for Red Hat OpenStack Platform 16.1

Last Updated: 2023-10-03

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

Abstract

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

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see our CTO Chris Wright's message.

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your input on our documentation. Tell us how we can make it better.

Using the Direct Documentation Feedback (DDF) function

Use the **Add Feedback** DDF function for direct comments on specific sentences, paragraphs, or code blocks.

1. View the documentation in the *Multi-page HTML* format.
2. Ensure that you see the **Feedback** button in the upper right corner of the document.
3. Highlight the part of text that you want to comment on.
4. Click **Add Feedback**.
5. Complete the **Add Feedback** field with your comments.
6. Optional: Add your email address so that the documentation team can contact you for clarification on your issue.
7. Click **Submit**.

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 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)

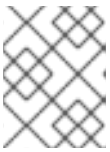

NOTE

Because Internet Explorer 11 is no longer maintained, expect a degradation of functionality when displaying the dashboard.

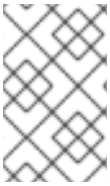

NOTE

Before you deploy 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). You can use Bare Metal Provisioning to provision bare-metal machines by 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 (RHOSP) 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.

NOTE

If you synchronize repositories by using Red Hat Satellite, you can enable specific versions of the Red Hat Enterprise Linux repositories. However, the repository label remains the same despite the version you choose. For example, if you enable the 8.2 version of the BaseOS repository, the repository name includes the specific version that you enabled, but the repository label is still **rhel-8-for-x86_64-baseos-tus-rpms**.

WARNING

Any repositories outside the ones specified here are not supported. Unless recommended, do not enable any other products or repositories outside the ones listed in the following tables or else you might encounter package dependency issues. Do not enable Extra Packages for Enterprise Linux (EPEL).

Core repositories

The following table lists core repositories for installing the undercloud.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8.2 for x86_64 - BaseOS (RPMs) Telecommunications Update Service (TUS)	**rhel-8-for-x86_64-baseos-tus-rpms**	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 8.2 for x86_64 - AppStream (RPMs)	**rhel-8-for-x86_64-appstream-tus-rpms**	Contains Red Hat OpenStack Platform dependencies.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8.2 for x86_64 - High Availability (RPMs) Telecommunications Update Service (TUS)	**rhel-8-for-x86_64-highavailability-tus-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.
Advanced Virtualization for RHEL 8 x86_64 (RPMs)	**advanced-virt-for-rhel-8-x86_64-eus-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.
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 or if you want to integrate with an existing Ceph Storage cluster.

IBM POWER repositories

The following table contains a list of repositories for RHOSP 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 Fast Datapath for RHEL 8 IBM Power, little endian (RPMS)	**fast-datapath-for-rhel-8-ppc64le-rpms**	Provides Open vSwitch (OVS) packages for OpenStack Platform.
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.
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 (RHOSP) 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.

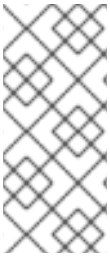

NOTE

If you synchronize repositories by using Red Hat Satellite, you can enable specific versions of the Red Hat Enterprise Linux repositories. However, the repository label remains the same despite the version you choose. For example, if you enable the 8.2 version of the BaseOS repository, the repository name includes the specific version that you enabled, but the repository label is still **rhel-8-for-x86_64-baseos-tus-rpms**.

WARNING

Any repositories outside the ones specified here are not supported. Unless recommended, do not enable any other products or repositories outside the ones listed in the following tables or else you might encounter package dependency issues. Do not enable Extra Packages for Enterprise Linux (EPEL).

Controller node repositories

The following table lists core repositories for Controller nodes in the overcloud.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8.2 for x86_64 - BaseOS (RPMs) Telecommunications Update Service (TUS)	**rhel-8-for-x86_64-baseos-tus-rpms**	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 8.2 for x86_64 - AppStream (RPMs)	**rhel-8-for-x86_64-appstream-tus-rpms**	Contains Red Hat OpenStack Platform dependencies.
Red Hat Enterprise Linux 8.2 for x86_64 - High Availability (RPMs) Telecommunications Update Service (TUS)	**rhel-8-for-x86_64-highavailability-tus-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-eus-rpms**	Provides virtualization packages for OpenStack Platform.
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.
Red Hat Ceph Storage Tools 4 for RHEL 8 x86_64 (RPMs)	**rhceph-4-tools-for-rhel-8-x86_64-rpms**	Tools for Red Hat Ceph Storage 4 for Red Hat Enterprise Linux 8.
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.

Compute and ComputeHCI node repositories

The following table lists core repositories for Compute and ComputeHCI nodes in the overcloud.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8.2 for x86_64 - BaseOS (RPMs) Telecommunications Update Service (TUS)	**rhel-8-for-x86_64-baseos-tus-rpms**	Base operating system repository for x86_64 systems.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8.2 for x86_64 - AppStream (RPMs)	**rhel-8-for-x86_64-appstream-tus-rpms**	Contains Red Hat OpenStack Platform dependencies.
Red Hat Enterprise Linux 8.2 for x86_64 - High Availability (RPMs) Telecommunications Update Service (TUS)	**rhel-8-for-x86_64-highavailability-tus-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-eus-rpms**	Provides virtualization packages for OpenStack Platform.
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.
Red Hat Ceph Storage Tools 4 for RHEL 8 x86_64 (RPMs)	**rhceph-4-tools-for-rhel-8-x86_64-rpms**	Tools for Red Hat Ceph Storage 4 for Red Hat Enterprise Linux 8.
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.

Real Time Compute 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.

Name	Repository	Description of requirement
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.

Ceph Storage node repositories

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

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8.2 for x86_64 - BaseOS (RPMs) Telecommunications Update Service (TUS)	**rhel-8-for-x86_64-baseos-tus-rpms**	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 8.2 for x86_64 - AppStream (RPMs)	**rhel-8-for-x86_64-appstream-tus-rpms**	Contains Red Hat OpenStack Platform dependencies.
Red Hat Enterprise Linux 8.2 for x86_64 - High Availability (RPMs) Telecommunications Update Service (TUS)	**rhel-8-for-x86_64-highavailability-tus-rpms**	High availability tools for Red Hat Enterprise Linux. NOTE: If you used the **overcloud-full** image for your Ceph Storage role, you must enable this repository. Ceph Storage roles should use the **overcloud-minimal** image, which does not require this repository.
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. This repository is included with standalone Ceph Storage subscriptions. If you use a combined OpenStack Platform and Ceph Storage subscription, use the **openstack-16.1-for-rhel-8-x86_64-rpms** repository.

Name	Repository	Description of requirement
Red Hat OpenStack Platform 16.1 for RHEL 8 (RPMs)	**openstack-16.1-for-rhel-8-x86_64-rpms**	Packages to help director configure Ceph Storage nodes. This repository is included with combined OpenStack Platform and Ceph Storage subscriptions. If you use a standalone Ceph Storage subscription, use the **openstack-16.1-deployment-tools-for-rhel-8-x86_64-rpms** repository.
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.
Red Hat Fast Datapath for RHEL 8 (RPMS)	**fast-datapath-for-rhel-8-x86_64-rpms**	Provides Open vSwitch (OVS) packages for OpenStack Platform. If you are using OVS on Ceph Storage nodes, add this repository to the network interface configuration (NIC) templates.

IBM POWER repositories

The following table lists repositories for RHOSP 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 Fast Datapath for RHEL 8 IBM Power, little endian (RPMS)	**fast-datapath-for-rhel-8-ppc64le-rpms**	Provides Open vSwitch (OVS) packages for OpenStack Platform.

Name	Repository	Description of requirement
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 through 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.

1.12. UNSUPPORTED FEATURES

The following features are not supported in Red Hat OpenStack Platform:

- Custom policies, which includes modification of **policy.json** files either manually or through any ***Policies** heat parameters. Do not modify the default policies unless the documentation contains explicit instructions to do so.
- Containers are not available for the following packages, therefore they are not supported in RHOSP:
  - **nova-serialproxy**

    - **nova-spicehtml5proxy**

- File injection of personality files to inject user data into virtual machine instances. Instead, cloud users can pass data to their instances by using the **--user-data** option to run a script during instance boot, or set instance metadata by using the **--property** option when launching an instance. For more information, see Creating a customized instance.

- Upgrades from RHOSP 15 to 16.1.

If you require support for any of these features, please contact the Red Hat Customer Experience and Engagement team to obtain a support exception.

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 (nova).

Using the Placement service for Tenant-isolated host aggregates

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 project-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. EDGE COMPUTING

This section outlines the top new features for edge computing.

Edge features added in Red Hat OpenStack Platform 16.1.2

Edge support is now available for Ansible-based transport layer security everywhere (TLSe), Key Manager service (barbican), and routed provider networks. You can now use an Ansible playbook to pre-cache Image service (glance) images for edge sites.

2.4. NETWORKING

This section outlines the top new features for the Networking service (neutron).

ML2/OVS supports QoS policies on hardware offloaded direct ports

Starting in Red Hat OpenStack Platform 16.1.7, the Modular Layer 2 plug-in with the Open vSwitch mechanism driver (ML2/OVS) now supports QoS rules on hardware offloaded *direct* ports.

New Networking service quota driver no longer requires MariaDB resource request locks

Starting in Red Hat OpenStack Platform (RHOSP) 16.1.7, there is a new quota driver for the RHOSP Networking service (neutron). The Networking service no longer uses global locks, but instead uses MariaDB transaction isolation levels to retrieve used resources and the current resource reservations. This new driver creates a reservation within the same database transaction, provided that the quota is not breached.

Because the new transaction both counts the used resources and creates the reservation, this new quota driver is much faster than the previous driver that uses the **make_reservation** transaction.

Using the new quota driver, the Networking service is less likely to encounter resource request bottlenecks that can lead to the database locking up.

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 active-standby topology for Load-balancing service instances in the Using Octavia for Load Balancing-as-a-Service 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 Using Octavia for Load Balancing-as-a-Service guide.

Routed provider networks

Starting in Red Hat OpenStack Platform 16.1.1, you can use the ML2/OVS or the SR-IOV mechanism drivers to deploy routed provider networks. Routed provider networks are common in edge distributed compute node (DCN) and spine-leaf routed data center deployments. Routed provider networks enable a single provider network to represent multiple layer 2 networks (broadcast domains) or network segments, which permits the operator to present only one network to users. For more information, see Deploying routed provider networks in the Networking Guide.

SR-IOV with native OVN DHCP in ML2/OVN deployments

Starting in Red Hat OpenStack Platform 16.1.1, you can use SR-IOV with native OVN DHCP (no need for neutron DHCP) in ML2/OVN deployments.

For more information, see Enabling SR-IOV with ML2/OVN and Native OVN DHCP and Limits of the ML2/OVN mechanism driver in the Networking Guide.

Northbound path MTU discovery support for jumbo frames

Red Hat OpenStack Platform 16.1.2 introduces MTU discovery to support UDP jumbo frames. After receiving a jumbo UDP frame that exceeds the MTU of the external network, ML2/OVN routers return ICMP "fragmentation needed" packets back to the sending VM. The sending application can then break the payload into smaller packets. Previously, the inability to return ICMP "fragmentation needed" packets resulted in packet loss. For more information about the necessary configuration steps, see Configuring ML2/OVN northbound path MTU discovery for jumbo frame fragmentation in the Advanced Overcloud Customization guide.

Note that in east/west traffic OVN does not support fragmentation of packets that are larger than the smallest MTU on the east/west path.

Example

- VM1 is on Network1 with an MTU of 1300.
- VM2 is on NEtwork2 with an MTU of 1200.
- A ping in either direction between VM1 and VM2 with a size of 1171 or less succeeds. A ping with a size greater than 1171 results in 100 percent packet loss.
  See https://bugzilla.redhat.com/show_bug.cgi?id=1891591.

Load-balancing service instance (amphora) log offloading

By default, Load-balancing service instances (amphorae) store logs on the local machine in the systemd journal. However, starting in Red Hat OpenStack Platform 16.1.2, you can specify that amphorae offload logs to syslog receivers to aggregate both administrative and tenant traffic flow

logs. Log offloading enables administrators to go to one location for logs, and retain logs when amphorae are rotated. For more information, see Basics of offloading Load-balancing service instance (amphora) logs in the Using Octavia for Load Balancing-as-a-Service guide.

OVN provider driver for the Load-balancing service (octavia)

Red Hat OpenStack Platform (RHOSP) 16.1.2 introduces full support for the Open Virtual Network (OVN) load-balancing provider, which is a lightweight load-balancer with a basic feature set. Typically used for east-west, layer 4 network traffic, OVN provisions fast and consumes less resources than a full-featured load-balancing provider such as amphora.

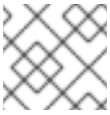

NOTE

Health check functionality is not implemented for the OVN provider driver.

On RHOSP deployments that use the ML2/OVN neutron plug-in, RHOSP director automatically enables the OVN provider driver in the Load-balancing service (octavia), without requiring additional installation or configuration steps. As with all RHOSP deployments, the default load-balancing provider driver, amphora, remains enabled and fully supported. For more information, see Creating an OVN load balancer in the Using Octavia for Load Balancing-as-a-Service guide.

In-place migration from ML2/OVS to ML2/OVN—supported in RHOSP 16.2

If your existing Red Hat OpenStack Platform (RHOSP) deployment uses the ML2/OVS mechanism driver, you must evaluate the benefits and feasibility of replacing the OVS driver with the ML2/OVN mechanism driver. Red Hat does not support a direct migration to ML2/OVN in RHOSP 16.1. You must upgrade to the latest RHOSP 16.2 version before migrating to the ML2/OVN mechanism driver.

NOTE

Red Hat requires that you file a preemptive support case before attempting a migration from ML2/OVS to ML2/OVN. Red Hat does not support migrations without the preemptive support case.

2.5. 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 File Systems service (manila) fully supports the Native CephFS driver.

FileStore to BlueStore OSD migration

Starting in Red Hat OpenStack Platform 16.1.2, an Ansible driven workflow migrates Ceph OSDs from FileStore to BlueStore. This means that customers who are using direct-deployed Ceph Storage can complete the Framework for Upgrades (OSP13 to OSP16.1) process.

In-use RBD volume migration

Starting in Red Hat OpenStack Platform 16.1.2, you can migrate or retype RBD in-use cinder volumes from one Ceph pool to another within the same Ceph cluster. See https://bugzilla.redhat.com/show_bug.cgi?id=1293440.

Red Hat OpenShift Container Platform support

The Shared File Systems service (manila) with CephFS through NFS fully supports serving shares to Red Hat OpenShift Container Platform through Manila CSI. This solution is not intended for large scale deployments. For important recommendations, see https://access.redhat.com/articles/6667651.

2.6. 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 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.

2.7. CLOUDOPS

This section outlines the top new features and changes for the CloudOps components.

Native multiple cloud support

In Service Telemetry Framework (STF) 1.1, multiple cloud support is native in the Service Telemetry Operator. This is provided by the new **clouds** parameter.

Custom SmartGateway objects

In STF 1.1, the Smart Gateway Operator can directly manage custom SmartGateway objects. You can use the **clouds** parameter to configure STF-managed cloud instances. You can set the **clouds** object to an empty set to indicate the Service Telemetry Operator does not manage SmartGateway objects.

SNMP traps

In STF 1.1, delivery of SNMP traps through Alertmanager webhooks has been implemented.

2.8. 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-coverged infrastructure (HCI) deployments with OVS-DPDK. In an HCI architecture, overcloud nodes with Compute and Ceph Storage services are colocated and configured for optimized resource usage.

Open vSwitch (OVS) hardware offload with OVS-ML2 or OVN-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. See Configuring OVS hardware offload in the Network Functions Virtualization Planning and Configuration Guide.

2.9. TECHNOLOGY PREVIEWS

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

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 . In RHOSP 16.1.2 this feature has graduated from Technology Preview to full support.

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.

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 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 the Block Storage service 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 OSP10-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 Block Storage service (cinder) and etcd services were configured with malformed etcd URIs, and the Block Storage service and etcd services failed on startup.
With this update, the IPv6 addresses in the etcd URI are correct and the Block Storage service 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, Block Storage service and Compute 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 Block Storage and Compute 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 correctly:

- If you specify a **source_volid** 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 Block Storage service sets the volume type as 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 the Bare Metal Provisioning service (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 two 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 Red Hat OpenStack Platform with Ceph Storage 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 (RHOSP) to use an external, pre-existing Ceph RadosGW cluster. You can manage this cluster externally as an object-store for RHOSP 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. See Configuring OVS hardware offload in the Network Functions Virtualization Planning and Configuration Guide.

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 by 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-coverged infrastructure (HCI) deployments with OVS-DPDK. In an HCI architecture, overcloud nodes with Compute and Ceph Storage services are colocated 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#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 DCHP 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 the Compute service (nova) scheduler is not segment-aware, you must map each leaf, rack segment, or DCN edge site to a Compute service host-aggregate or availability zone. If the deployments require DHCP or the metadata service, you must also define a Compute service availability zone for each edge site or segment.

Known limitations:

- Supported with ML2/OVS only. Not supported with ML2/OVN (RFE Bug 1797664).
- Compute service scheduler is not segment-aware. For successful Compute service scheduling, map each segment or edge site to a Compute service host-aggregate or availability zone. Currently there are only two instance boot options available [RFE Bug 1761903]:
  - Boot an instance with port-id and using a different 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 you use 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 int he **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-ansible 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 File Systems 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 Systems 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 Systems 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 Systems service with CephFS through 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#1837316

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 <heath_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.

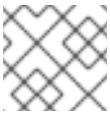

NOTE

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.

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**.
  Workaround: nbind or detach VFs before you perform these actions: . Shut down and detach any VMs. . Remove VF LAG BOND interface from OVS. . Unbind each configured VF: **# echo <VF PCIe BDF> > /sys/bus/pci/drivers/mlx5_core/unbind** . 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 a 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.*
```

BZ#1980829

If you change the **TRIPLEO_HEAT_TEMPLATE_KERNEL_ARGS**, such as the value for the **hugepages** parameter, during a fast-forward upgrade (FFU) from Red Hat OpenStack Platform (RHOSP) 13 to RHOSP 16.1, the upgrade fails because of duplicate entries for the kernel args. Avoid changing the kernel args during FFU.

Workaround: In RHOSP 16.1 you can manually change the kernel args, which are usually located at /usr/share/ansible/roles/tripleo-kernel/tasks/kernelargs.yml

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.

3.2. RED HAT OPENSTACK PLATFORM 16.1.1 MAINTENANCE RELEASE - AUGUST 27, 2020

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.

3.2.1. Bug fix

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

BZ#1845726

This director enhancement automatically installs the Leapp utility on overcloud nodes to prepare for OpenStack upgrades.https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.1/html-single/release_notes/index This enhancement includes two new Heat parameters: LeappRepoInitCommand and LeappInitCommand. In addition, if you have the following repository defaults, you do not need to pass UpgradeLeappCommandOptions values.

> --enablerepo rhel-8-for-x86_64-baseos-eus-rpms --enablerepo rhel-8-for-x86_64-appstream-eus-rpms --enablerepo rhel-8-for-x86_64-highavailability-eus-rpm1866372s --enablerepo advanced-virt-for-rhel-8-x86_64-rpms --enablerepo ansible-2.9-for-rhel-8-x86_64-rpms --enablerepo fast-datapath-for-rhel-8-x86_64-rpms

BZ#1847463

This update fixes a bug that caused the **generate-inventory** step to fail during in-place migration from ML2/OVS to ML2/OVN.
Note that in the Red Hat OpenStack Platform 16.1.0 (GA release), migration from ML2/OVS to ML2/OVN was not supported. As of Red Hat OpenStack Platform 16.1.1, in-place migration is supported for non-NFV deployments, with various exceptions, limitations, and requirements as described in "Migrating from ML2/OVS to ML2/OVN." [1]

[1] https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.1/html-single/networking_with_open_virtual_network/index#migrating-ml2ovs-to-ovn

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 Red Hat OpenStack Platform (RHOSP) 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#1851914

The overcloud deployment steps included an older Ansible syntax that tagged the **tripleo-bootstrap** and **tripleo-ssh-known-hosts** roles as **common_roles**. This older syntax caused Ansible to run tasks tagged with the **common_roles** when Ansible did not use the **common_roles** tag. This syntax resulted in errors during the 13 to 16.1 **system_upgrade** process.
This update uses a newer syntax to tag the **tripleo-bootstrap** and **tripleo-ssh-known-hosts** roles as **common_roles**. Errors do not appear during the 13 to 16.1 **system_upgrade** process and you no longer include the **--playbook upgrade_steps_playbook.yaml** option to the **system_upgrade** process as a workaround.

BZ#1852620

This update fixes a bug that prevented the successful deployment of transport layer security (TLS) everywhere with public TLS certifications.

BZ#1852868

This update fixes a Red Hat Ceph Storage (RHCS) version compatibility issue that caused failures during upgrades from Red Hat OpenStack platform 13 to 16.1. Before this fix, validations performed during the upgrade worked with RHCS3 clusters but not RHCS4 clusters. Now the validation works with both RHCS3 and RHCS4 clusters.

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 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#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.

BZ#1858673

This update fixes a GRUB parameter naming convention that led to unpredictable behaviors on compute nodes during leapp upgrades.
Previously, the presence of the obsolete "TRIPELO" prefix on GRUB parameters caused problems.

The file /etc/default/grub has been updated with GRUB for the tripleo kernel args parameter so that leapp can upgrade it correctly. This is done by adding "upgrade_tasks" to the service "OS::TripleO::Services::BootParams", which is a new service added to all roles in the roles_data.yaml file.

BZ#1866372

This update fixes a problem that caused baremetal nodes to become non-responsive during Leapp upgrades.
Previously, Leapp did not process transient interfaces like SR-IOV virtual functions (VF) during migration. As a result, Leapp did not find the VF interfaces during the upgrade, and nodes entered an unrecoverable state.

Now the service "OS::TripleO::Services::NeutronSriovAgent" sets the physical function (PF) to remove all VFs, and migrates workloads before the upgrade. After the successful Leapp upgrade, os-net-config runs again with the "--no-activate" flag to re-establish the VFs.

3.2.2. Enhancements

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

BZ#1666684

In this release, you can use SR-IOV in an ML2/OVN deployment with native OVN DHCP. SR-IOV in an ML2/OVN deployment no longer requires the Networking service (neutron) DHCP agent.
When virtual machines boot on hypervisors that support SR-IOV NICs, the OVN controllers on the controller or network nodes can reply to the DHCP, internal DNS, and IPv6 router solicitation requests from the virtual machine.

This feature was available as a technology preview in RHOSP 16.1.0. Now it is a supported feature.

The following limitations apply to the feature in this release:

- All external ports are scheduled on a single gateway node because there is only one HA Chassis Group for all of the ports.
- North/south routing on VF(direct) ports on VLAN tenant networks does not work with SR-IOV because the external ports are not colocated with the logical router's gateway ports. See https://bugs.launchpad.net/neutron/+bug/1875852.

BZ#1671811

In the first maintenance release of Red Hat OpenStack Platform 16.1 there is support for routed provider networks using the ML2/OVS and SR-IOV mechanism drivers.
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 and spine-leaf routed data center deployments.

For more information, see https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.1/html-single/networking_guide/index#deploy-routed-prov-networks_rhosp-network.

3.2.3. 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#1801721

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

3.2.4. Known issues

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

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#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#1866562

Currently, you cannot scale down or delete compute nodes if Red Hat OpenStack Platform is deployed with TLS-e using tripleo-ipa. This is because the cleanup role, traditionally delegated to the undercloud as localhost, is now being invoked from the mistral container.
For more information, see https://access.redhat.com/solutions/5336241

BZ#1867458

A Leapp issue causes failure of fast forward upgrades from Red Hat OpenStack (RHOSP) platform 13 to RHOSP 16.1.
A Leapp upgrade from RHEL 7 to RHEL 8 removes all older RHOSP packages and performs an operating system upgrade and reboot. Because Leapp installs os-net-config package at the "overcloud upgrade run" stage, os-net-config-sriov executable is not available for sriov_config serivce to configure virtual functions (VF) and switchdev mode after reboot. As a result, VFs are not configured and switchdevmode is not applied on the physical function (PF) interfaces.

As a workaround, manually create the VFs, apply switchdevmode to the VF interface, and restart the VF interface.

3.3. RED HAT OPENSTACK PLATFORM 16.1.2 MAINTENANCE RELEASE - OCTOBER 27, 2020

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.

3.3.1. Bug fix

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

BZ#1721361

This update includes the following bug fix patches related to fully qualified domain names (FQDN).

- *Kaminario Fix unique_fqdn_network option*
  Previously, the Kaminario driver accepted the unique_fqdn_network configuration option in the specific driver section. When this option was moved, a regression was introduced: the parameter was now only used if it was defined in the shared configuration group.

  This patch fixes the regression and makes it possible to define the option in the shared configuration group as well as the driver specific section.

- *HPE 3PAR Support duplicated FQDN in network*

The 3PAR driver uses the FQDN of the node that is doing the attach as an unique identifier to map the volume.

Because the FQDN is not always unique, in some environments the same FQDN can be found in different systems. In those cases, if both try to attach volumes, the second system will fail.

For example, this could happen in a QA environment where VMs share names like controller-.localdomain and compute-0.localdomain.

This patch adds the **unique_fqdn_network** configuration option to the 3PAR driver to prevent failures caused by name duplication between systems. (BZ#1721361)

BZ#1792500

Inadequate timeout values can cause an overcloud deployment to fail after four hours. To prevent these timeout failures, set the following undercloud and overcloud timeout parameters:

- Undercloud timeouts (seconds):

  Example

  ```
  parameter_defaults:
    TokenExpiration: 86400
    ZaqarWsTimeout: 86400
  ```

- Overcloud deploy timeouts (minutes):

  Example

  ```
  $ openstack overcloud deploy --timeout 1440
  ```

  The timeouts are now set.

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_volid** 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 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#1843789

Before this update, when you created a volume from a snapshot, the operation could fail because the Block Storage service (cinder) would try to assign the default volume type to the new volume instead of inferring the correct volume type from the snapshot. With this update, you no longer have to specify the volume type when you create a volume.

BZ#1848420

This update makes it possible to run the Brocade FCZM driver in RHOSP 16.
The Brocade FCZM vendor chose not to update the driver for Python 3, and discontinued support of the driver past the Train release of OpenStack [1]. Red Hat OpenStack (RHOSP) 16 uses Python 3.6.

The upstream Cinder community assumed the maintenance of the Brocade FCZM driver on a best-effort basis, and the bugs that prevented the Brocade FCZM from running in a Python 3 environment (and hence in RHOSP 16) have been fixed.

[1] https://docs.broadcom.com/doc/12397527

BZ#1855112

This update increases the speed of stack updates in certain cases.
Previously, stack update performance was degraded when the Ansible --limit option was not passed to ceph-ansible. During a stack update, ceph-ansible sometimes made idempotent updates on nodes even if the --limit argument was used.

Now director intercepts the Ansible --limit option and passes it to the ceph-ansible excecution. The --limit option passed to commands starting with 'openstack overcloud' deploy is passed to the ceph-ansible execution to reduce the time required for stack updates.

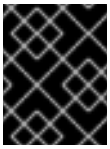

IMPORTANT

Always include the undercloud in the limit list when using this feature with ceph-ansible.

BZ#1855751

Before this update, to successfully run a leapp upgrade during the Framework for Upgrades upgrade (FFU) from RHOSP 13 to RHOSP 16.1, the node where the Red Hat Enterprise Linux upgrade was occurring had to have the **PermitRootLogin** field defined in the ssh config file (**/etc/ssh/sshd_config**).
With this update, the Orchestration service (heat) no longer requires you to modify **/etc/ssh/sshd_config** with the **PermitRootLogin** field.

BZ#1862213

This update fixes a problem that caused volume attachments to fail on a VxFlexOS cinder backend.
Previously, attempts to attach a volume on a VxFlexOS cinder backend failed because the cinder driver for the VxFlexOS back end did not include all of the information required to connect to the volume.

The VxFlexOS cinder driver has been updated to include all the information required in order to connect to a volume. The attachments now work correctly.

BZ#1868620

This update fixes incorrect parameter names in Dell EMC Storage Templates.

BZ#1869346

This update fixes an incompatibility that caused VxFlex volume detachment attempts to fail.

A recent change ix n VxFlex cinder volume credentialing methods was not backward compatible with pre-existing volume attachments. If a VxFlex volume attachment was made before the credentialing method change, attempts to detach the volume failed.

Now the detachments do not fail.

BZ#1872211

This update modifies **get_device_info** to use lsscsi to get **[H:C:T:L]** values, making it possible to support more than 255 logical unit numbers (LUNs) and host logical unit (HLU) ID values. Previously, **get_device_info** used sg_scan to get these values, with a limit of 255.

You can get two device types with **get_device_info**:

- o /dev/disk/by-path/xxx, which is a symlink to /dev/sdX
- o /dev/sdX
  sg_scan can process any device name, but lsscsi only shows /dev/sdx names.

  If the device is a symlink, **get_device_info** uses the device name that the device links to. Otherwise **get_device_info** uses the device name directly.

  Then **get_devix ce_info** gets the device info '[H:C:T:L]' by comparing the device name with the last column of lsscsi output.

BZ#1873329

This update fixes a bug that prevented the distributed compute nodes (DCN) compute servcie from accessing the glance service.
Previously, distributed compute nodes were configured with a glance endpoint URI that specified an IP address, even when deployed with internal transport layer security (TLS). Because TLS requires the endpoint URI to specify a fully qualified domain name (FQDN), the compute service could not access the glance service.

Now, when deployed with internal TLS, DCN services are configured with glance endpoint URI that specifies a FQDN, and the DCN compute service can access the glance service.

BZ#1879190

This bug fix enables you to boot an instance from an encrypted volume when that volume was created from an image that in turn was created by uploading an encrypted volume to the Image Service as an image.==== Bug Fix

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

BZ#1721361

This update includes the following bug fix patches related to fully qualified domain names (FQDN).

- *Kaminario Fix unique_fqdn_network option*
  Previously, the Kaminario driver accepted the unique_fqdn_network configuration option in the specific driver section. When this option was moved, a regression was introduced: the parameter was now only used if it was defined in the shared configuration group.

  This patch fixes the regression and makes it possible to define the option in the shared configuration group as well as the driver specific section.

- *HPE 3PAR Support duplicated FQDN in network*

The 3PAR driver uses the FQDN of the node that is doing the attach as an unique identifier to map the volume.

Because the FQDN is not always unique, in some environments the same FQDN can be found in different systems. In those cases, if both try to attach volumes, the second system will fail.

For example, this could happen in a QA environment where VMs share names like controller-.localdomain and compute-0.localdomain.

This patch adds the **unique_fqdn_network** configuration option to the 3PAR driver to prevent failures caused by name duplication between systems. (BZ#1721361)

BZ#1792500

Inadequate timeout values can cause an overcloud deployment to fail after four hours. To prevent these timeout failures, set the following undercloud and overcloud timeout parameters:

- Undercloud timeouts (seconds):

  Example

  ```
  parameter_defaults:
    TokenExpiration: 86400
    ZaqarWsTimeout: 86400
  ```

- Overcloud deploy timeouts (minutes):

  Example

  ```
  $ openstack overcloud deploy --timeout 1440
  ```

  The timeouts are now set.

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_volid** 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 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#1843789

Before this update, when you created a volume from a snapshot, the operation could fail because the Block Storage service (cinder) would try to assign the default volume type to the new volume instead of inferring the correct volume type from the snapshot. With this update, you no longer have to specify the volume type when you create a volume.

BZ#1848420

This update makes it possible to run the Brocade FCZM driver in RHOSP 16.
The Brocade FCZM vendor chose not to update the driver for Python 3, and discontinued support of the driver past the Train release of OpenStack [1]. Red Hat OpenStack (RHOSP) 16 uses Python 3.6.

The upstream Cinder community assumed the maintenance of the Brocade FCZM driver on a best-effort basis, and the bugs that prevented the Brocade FCZM from running in a Python 3 environment (and hence in RHOSP 16) have been fixed.

[1] https://docs.broadcom.com/doc/12397527

BZ#1855112

This update increases the speed of stack updates in certain cases.
Previously, stack update performance was degraded when the Ansible --limit option was not passed to ceph-ansible. During a stack update, ceph-ansible sometimes made idempotent updates on nodes even if the --limit argument was used.

Now director intercepts the Ansible --limit option and passes it to the ceph-ansible excecution. The --limit option passed to commands starting with 'openstack overcloud' deploy is passed to the ceph-ansible execution to reduce the time required for stack updates.

IMPORTANT

Always include the undercloud in the limit list when using this feature with ceph-ansible.

BZ#1855751

Before this update, to successfully run a leapp upgrade during the Framework for Upgrades upgrade (FFU) from RHOSP 13 to RHOSP 16.1, the node where the Red Hat Enterprise Linux upgrade was occurring had to have the **PermitRootLogin** field defined in the ssh config file (**/etc/ssh/sshd_config**).
With this update, the Orchestration service (heat) no longer requires you to modify **/etc/ssh/sshd_config** with the **PermitRootLogin** field.

BZ#1862213

This update fixes a problem that caused volume attachments to fail on a VxFlexOS cinder backend.
Previously, attempts to attach a volume on a VxFlexOS cinder backend failed because the cinder driver for the VxFlexOS back end did not include all of the information required to connect to the volume.

The VxFlexOS cinder driver has been updated to include all the information required in order to connect to a volume. The attachments now work correctly.

BZ#1868620

This update fixes incorrect parameter names in Dell EMC Storage Templates.

BZ#1869346

This update fixes an incompatibility that caused VxFlex volume detachment attempts to fail.

A recent change ix n VxFlex cinder volume credentialing methods was not backward compatible with pre-existing volume attachments. If a VxFlex volume attachment was made before the credentialing method change, attempts to detach the volume failed.

Now the detachments do not fail.

BZ#1872211

This update modifies **get_device_info** to use lsscsi to get **[H:C:T:L]** values, making it possible to support more than 255 logical unit numbers (LUNs) and host logical unit (HLU) ID values. Previously, **get_device_info** used sg_scan to get these values, with a limit of 255.

You can get two device types with **get_device_info**:

- o /dev/disk/by-path/xxx, which is a symlink to /dev/sdX
- o /dev/sdX
  sg_scan can process any device name, but lsscsi only shows /dev/sdx names.

  If the device is a symlink, **get_device_info** uses the device name that the device links to. Otherwise **get_device_info** uses the device name directly.

  Then **get_devix ce_info** gets the device info '[H:C:T:L]' by comparing the device name with the last column of lsscsi output.

BZ#1873329

This update fixes a bug that prevented the distributed compute nodes (DCN) compute servcie from accessing the glance service.
Previously, distributed compute nodes were configured with a glance endpoint URI that specified an IP address, even when deployed with internal transport layer security (TLS). Because TLS requires the endpoint URI to specify a fully qualified domain name (FQDN), the compute service could not access the glance service.

Now, when deployed with internal TLS, DCN services are configured with glance endpoint URI that specifies a FQDN, and the DCN compute service can access the glance service.

BZ#1879190

This bug fix enables you to boot an instance from an encrypted volume when that volume was created from an image that in turn was created by uploading an encrypted volume to the Image Service as an image.

3.3.2. Enhancements

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

BZ#1293440

This update enables you to migrate or retype RBD in-use cinder volumes from one Ceph pool to another within the same Ceph cluster. For more information, see Basic volume usage and configuration in the Storage Guide

BZ#1628811

This update adds NIC partitioning support on Intel and Mellanox NICs.

BZ#1668213

This update introduces support for encrypted images with keys managed by the Key Manager service (barbican).
For some secure workflows in which at-rest data must remain encrypted, you can upload carefully prepared encrypted images into the Image service (glance) for consumption by the Block Storage service (cinder).

BZ#1676631

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

BZ#1845422

When using multiple stores in the Image Service (glance), the image owner can delete an image copy from a specific store. In Red Hat OpenStack Platform 16.1.2, this feature moves from Technology Preview to full support.

BZ#1852851

This update adds support for encrypted volumes and images on distributed compute nodes (DCN). DCN nodes can now access the Key Manager service (barbican) running in the central control plane.

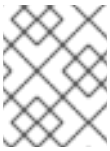

NOTE

This feature adds a new Key Manager client service to all DCN roles. To implement the feature, regenerate the **roles.yaml** file used for the DCN site's deployment.

For example:

```
$ openstack overcloud roles generate DistributedComputeHCI DistributedComputeHCIScaleOut -o ~/dcn0/roles_data.yaml
```

Use the appropriate path to the roles data file.

BZ#1859750

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#1859757

Previously, it was not possible to upgrade to TLS Everywhere in an existing deployment. With this update, you can secure the in-flight connections between internal OpenStack services without reinstallation.

BZ#1859758

You can use Atos Hardware Security Module (HSM) appliances in high availability (HA) mode with the Key Manager service (barbican). In Red Hat OpenStack Platform 16.1.2, this feature moves from Technology Preview to full support.

BZ#1862545

This release adds support for the Dell EMC PowerStore driver for the Block Storage service (cinder) back end.

BZ#1862546

This enhancement adds a new driver for the Dell EMC PowerStore to support Block Storage service back end servers.

BZ#1862547

This enhancement adds a new driver for the Dell EMC PowerStore to support Block Storage service back end servers.

BZ#1874847

This update introduces support of TLS Everywhere with Triple IPA for Distributed Compute Nodes (DCN).

BZ#1874863

The update introduces support of Networking service (neutron) routed provider networks with Distributed Compute Nodes (DCN).

BZ#1459187

Red Hat OpenStack Platform (RHOSP) 16.1 includes support 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. In RHOSP 16.1.2 this feature has graduated from Technology Preview to full support.

BZ#1474394

Red Hat OpenStack Platform (RHOSP) 16.1 includes support for bare metal provisioning over an IPv6 provisioning network for BMaaS (Bare Metal as-a-Service) tenants. In RHOSP 16.1.2, this feature has graduated from Technology Preview to full support.

3.3.3. 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#1703958

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

BZ#1820742

RHOSP 16.1.2 introduces a technology preview of the AMD EPYC 2 (Rome) platform with the UEFI setting **NPS** (Numa Per Socket)` set to **1**.
Other values of **NPS** (2 or 4) are used in DPDK benchmarks to reach the platform peak performances, without OpenStack, on bare metal.

Red Hat continues to evaluate the operational trade-off of **NPS=2** or **NPS=4** with OpenStack. This configuration exposes multiple Numa nodes per socket.

BZ#1827283

Red Hat OpenStack Platform 16.1.2 introduces a technology preview of the AMD EPYC 2 (Rome) platform with the UEFI setting **NPS** (Numa Per Socket) set to **1**.
Other values of **NPS** (2 or 4) are used in DPDK benchmarks to reach the platform peak performances, without OpenStack, on bare metal.

Red Hat continues to evaluate the operational trade-off of **NPS=2** or **NPS=4** with OpenStack. This configuration exposes multiple Numa nodes per socket.

BZ#1875310

Red Hat OpenStack Platform 16.1.2 introduces a technology preview of OVN and OVS-DPDK colocated with SR-IOV on the same hypervisor.
For related issues, see:

https://bugzilla.redhat.com/show_bug.cgi?id=1575512 and

https://bugzilla.redhat.com/show_bug.cgi?id=1575512

BZ#1875323

Red Hat OpenStack Platform 16.1.2 introduces a technology preview of OVN with OVS TC Flower-based offloads.
Note that VXLAN is not supported by OVN for regular inter-chassis communication. Thus, VXLAN with Hardware offload using OVN is not supported. See https://bugzilla.redhat.com/show_bug.cgi?id=1881704.

3.3.4. 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#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#1852082

In this update, the Red Hat OpenStack Platform (RHOSP) Orchestration service (heat) now enables you to deploy multiple Dell EMC XtremIO back ends with any combination of storage protocols for the Block Storage service (cinder).
A new heat parameter, **CinderXtremioStorageProtocol**, now enables you to choose between Fibre Channel (FC) or iSCSI storage protocols.

A new heat template enables you to deploy more than one XtremIO back end.

Previously, RHOSP director only supported one iSCSI back end for the Block Storage service. (The legacy iSCSI-only heat template will be deprecated in a future RHOSP release).

BZ#1852084

Red Hat OpenStack Platform 16.1.2 includes Orchestration service (heat) template support for the VXFlexOS driver for Block Storage service (cinder) back ends.

BZ#1852087

Red Hat OpenStack Platform 16.1.2 includes support for Dell EMC Storage Center (SC) back ends for the Block Storage service (cinder). The SC back end driver now supports both iSCSI and FC protocols, 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#1852088

PowerMax configuration options have changed after Red Hat OpenStack Platform 10 (newton). This update includes the latest PowerMax configuration options and supports both iSCSI and FC protocols.
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#1853450

Red Hat OpenStack Platform 16.1.2 includes Puppet support (**puppet-cinder** module) for the VXFlexOS driver for Block Storage service (cinder) back ends.

BZ#1853454

Red Hat OpenStack Platform 16.1.2 includes Puppet support (**puppet-tripleo** module) for the VXFlexOS driver for Block Storage service (cinder) back ends.

BZ#1877688

This update safeguards against potential package content conflict after content was moved from **openstack-tripleo-validations** to another package.

3.3.5. Known issues

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

BZ#1547074

Transmission of jumbo UDP frames on ML2/OVN routers depends on a kernel release that is not yet avaialbe.
After receiving a jumbo UDP frame that exceeds the maximum transmission unit of the external network, ML2/OVN routers can return ICMP "fragmentation needed" packets back to the sending VM, where the sending application can break the payload into smaller packets. To determine the packet size, this feature depends on discovery of MTU limits along the south-to-north path.

South-to-north path MTU discovery requires kernel-4.18.0-193.20.1.el8_2, which is scheduled for availability in a future release. To track availability of the kernel version, see https://bugzilla.redhat.com/show_bug.cgi?id=1860169.

BZ#1623977

When you enable Load-balancing service instance (amphora) log offloading, both the administrative logs and the tenant logs are written to the same file (**octavia-amphora.log**). This is a known issue caused by an incorrect default value for the Orchestration service (heat) parameter, **OctaviaTenantLogFacility**. As a workaround, perform the following steps:
Set **OctaviaTenantLogFacility** to zero (0) in a custom environment file and run the **openstack overcloud deploy** command:

```
parameter_defaults:
  OctaviaLogOffload: true
  OctaviaTenantLogFacility: 0
  ...
```

For more information, see Modifying the overcloud environment

BZ#1733577

A known issue causes the migration of Ceph OSDs from FileStore to BlueStore to fail. In use cases where the **osd_objectstore** parameter was not set explicitly when you deployed Red Hat OpenStack

Platform 13 with Red Hat Ceph Storage 3, the migration exits without converting any OSDs and falsely reports that the OSDs are already using BlueStore. For more information about the known issue, see https://bugzilla.redhat.com/show_bug.cgi?id=1875777
As a workaround, perform the following steps:

1. Include the following content in an environment file:

```
parameter_defaults:
  CephAnsibleExtraConfig:
    osd_objectstore: filestore
```

2. Perform a stack update with the **overcloud deploy --stack-only** command, and include the new or existing environment file that contains the **osd_objectstore** parameter. In the following example, this environment file is **<osd_objectstore_environment_file>**. Also include any other environment files that you included during the converge step of the upgrade:

```
$ openstack overcloud deploy --stack-only \
  -e <osd_objectstore_environment_file> \
  -e <converge_step_environment_files>
```

3. Proceed with the FileStore to BlueStore migration by using the existing documentation. See https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.1/html/framework_for_upgrades_13_to_16.1/OSD-migration-from-filestore-to-bluestore
Result: The FileStore to BlueStore playbook triggers the conversion process, and removes and re-creates the OSDs successfully.

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#1837316

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 <heath_monitor_id>**

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

BZ#1848462

Currently, on ML2/OVS and distributed virtual router (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 Red Hat OpenStack Platform 16.1 until this issue is fixed.

BZ#1861370

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.
As a 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.*
```

BZ#1866562

Currently, you cannot scale down or delete compute nodes if Red Hat OpenStack Platform is deployed with TLS Everywhere using tripleo-ipa. This is because the cleanup role, traditionally delegated to the undercloud as localhost, is now being invoked from the Workflow service (mistral) container.
For more information, see https://access.redhat.com/solutions/5336241

3.4. RED HAT OPENSTACK PLATFORM 16.1.3 MAINTENANCE RELEASE - DECEMBER 15, 2020

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.

3.4.1. Advisory list

This release includes the following advisories:

RHSA-2020:5411

Moderate: python-django-horizon security update

RHSA-2020:5412

Moderate: python-XStatic-jQuery224 security update

RHEA-2020:5413

Red Hat OpenStack Platform 16.1.3 bug fix and enhancement advisory

RHEA-2020:5414

Red Hat OpenStack Platform 16.1.3 director images bug fix advisory

RHEA-2020:5415

Red Hat OpenStack Platform 16.1.3 containers bug fix advisory

3.4.2. Bug fix

This bug was fixed in this release of Red Hat OpenStack Platform:

BZ#1878492

Before this update, director maintained Identity service (keystone) catalog entries for Block Storage service's (cinder) deprecated v1 API volume service, and the legacy Identity service endpoints were not compatible with recent enhancements to director's endpoint validations. As a result, stack updates failed if a legacy volume service was present in the Identity service catalog. With this update, director automatically removes the legacy volume service and its associated endpoints. Stack updates no longer fail Identity service endpoint validation.

3.4.3. Enhancements

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

BZ#1808577

This update supports the creation of volumes with tiering policy. There are four supported values:

- **StartHighThenAuto**(default)
- **Auto**
- **HighestAvailable**
- **LowestAvailable**

BZ#1862541

This enhancement adds a new driver for the Dell EMC PowerStore to support Block Storage service back end servers. The new driver supports the FC and iSCSI protocols, and includes these features:

- Volume create and delete
- Volume attach and detach
- Snapshot create and delete
- Create volume from snapshot
- Get statistics on volumes
- Copy images to volumes
- Copy volumes to images
- Clone volumes
- Extend volumes
- Revert volumes to snapshots

BZ#1809930

With this enhancement, the **OvsDpdkCoreList** parameter is now optional. If you set **OvsDpdkCoreList**, you pin the **ovs-vswitchd** non-pmd threads to the first core that you list in the parameter. If you exclude **OvsDpdkCoreList**, you enable the **ovs-vswitchd** non-pmd threads to use any non-isolated cores.

3.4.4. 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#1856404

In this release, the **collectd-libpod-stats** plugin collects CPU and memory metrics for containers running in the overcloud.

BZ#1867222

With this release, the VxFlex OS driver is renamed to PowerFlex. Names of configuration options have been changed and removed. The **ScaleIO** name and related **sio_** configuration options have been deprecated.

BZ#1867225

In this release, VxFlex OS driver is rebranded to PowerFlex.

3.4.5. Known issues

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

BZ#1261083

Currently, LVM filter is not set unless at least one device is listed in the **LVMFilterAllowlist** parameter.
Workaround: Set the **LVMFilterAllowdisk** parameter to contain at least one device, for example, the root disk. The LVM filter is set in **/etc/lvm/lvm.conf**.

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 pre-deployed Controller nodes:

```
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#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 a ceph-ansible bug. For more information, see https://bugzilla.redhat.com/show_bug.cgi?id=1860815.

BZ#1879418

It is a known issue that the **openstack overcloud status** command might not return the correct status for a given stack name when multiple stacks exist. Instead, the status of the most recently deployed stack is always returned, regardless of the stack name. This can lead to failure reported for all stacks when it is only the most recently deployed stack that has failed. Workaround: The true status of the deployment must be clear. For example, **openstack stack list** shows any overcloud deployment failures in the heat stage and the ansible deployment logs show failures in the config download stage.

BZ#1880979

Currently, a change in OSP13 puppet module kmod has resulted in the wrong module setting for **systemd-modules-load.service**. This is not an issue in OSP13 but results in failure during deployment in fast forward upgrade on OSP16.1.
Workaround: Enter the following command:

```
rm -f /etc/modules-load.d/nf_conntrack_proto_sctp.conf
```

BZ#1789822

Replacement of an overcloud Controller might cause swift rings to become inconsistent across nodes. This results in decreased availability of Object Storage service.
Workaround: Log in to the previously existing Controller node using SSH, deploy the updated rings, and restart the Object Storage containers:

```
(undercloud) [stack@undercloud-0 ~]$ source stackrc
(undercloud) [stack@undercloud-0 ~]$ nova list
...
| 3fab687e-99c2-4e66-805f-3106fb41d868 | controller-1 | ACTIVE | -          | Running     |
ctlplane=192.168.24.17 |
| a87276ea-8682-4f27-9426-6b272955b486 | controller-2 | ACTIVE | -          | Running     |
ctlplane=192.168.24.38 |
| a000b156-9adc-4d37-8169-c1af7800788b | controller-3 | ACTIVE | -          | Running     |
ctlplane=192.168.24.35
+
(undercloud) [stack@undercloud-0 ~]$ for ip in 192.168.24.17 192.168.24.38 192.168.24.35; do ssh
$ip 'sudo podman restart swift_copy_rings ; sudo podman restart $(sudo podman ps -a --format="
{{.Names}}" --filter="name=swift_*")'; done
```

BZ#1895887

After upgrading with the Leapp utility, Compute with OVS-DPDK workload does not function properly. To workaround this issue, perform one of the following steps:

- Remove the **/etc/modules-load.d/vfio-pci.conf** file before Compute upgrade.

or

- Restart **ovs-vswitchd** service on the Compute node after upgrade.

3.5. RED HAT OPENSTACK PLATFORM 16.1.4 MAINTENANCE RELEASE - MARCH 17, 2021

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.

3.5.1. Advisory list

This release includes the following advisories:

RHSA-2021:0915

Moderate: Red Hat OpenStack Platform 16.1.4 (python-django) security update

RHSA-2021:0916

Moderate: Red Hat OpenStack Platform 16.1.4 (etcd) security update

RHBA-2021:0817

Red Hat OpenStack Platform 16.1.4 director bug fix advisory

RHEA-2021:0918

Red Hat OpenStack Platform 16.1.4 director images bug fix advisory

RHEA-2021:0919

Red Hat OpenStack Platform 16.1.4 containers bug fix advisory

3.5.2. Bug fix

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

BZ#1728142

Before this update, the **--server** option was being ignored when passed with the **cinder service-get-log** command, which resulted in the logs for all hosts being returned instead of just the logs for a specific host. With this update, using the **--server** option correctly filters the logs for the specified host.

BZ#1828889

Before this update, the OVN mechanism driver did not correctly merge its agent list with those stored in the Networking (neutron) service database. With this update, the results from the OVN and Networking service database are merged before the API returns the result.

BZ#1847907

The 'all_tenants' key passed with a volume transfer request is removed because the database is unable to parse it. Removing this key allows the user to show the detail of a specific volume transfer by using the transfer name. Before this update, the 'all_tenants' key was removed only for admin users, which meant that non-admin users were unable to show volume transfers by using the transfer name. With this update, the 'all_tenants' key is now also removed for non-admins, allowing non-admins to show volume transfers by using the transfer name.

BZ#1874936

Before this update, TLS-E on pre-provisioned nodes failed with the message: "--server cannot be used without providing --domain". With this update, the IDM domain name is detected by first resolving "ipa-ca" through DNS, then doing a reverse DNS lookup on the resulting IP address. It might be necessary to add the PTR record, which is required for the reverse lookup, manually.

BZ#1881476

Before this update, if a user configured the **ContainerImagePrepare** parameter to use a custom tag, such as 'tag: "latest"' or 'tag: "16.1"', instead of the standard 'tag_from_label: "{version}-{release}"', the containers did not update to the latest container images.

With this update, the container images are always fetched anytime a user runs a deployment action, including updates, and the image ID is checked against the running container to see if it needs to be rebuilt to consume the latest image. Containers are now always refreshed during deployment actions and restarted if they are updated.

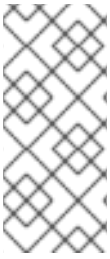

NOTE

This is a change from previous versions where the deployment checked only that the image existed rather than always fetching the image. If a user is reusing tags, for example, "latest", the containers might be updated on nodes if you perform actions such as scaling out. It is not recommended to use "latest" unless you are controlling container tags by using a Satellite server deployment.

BZ#1884556

Before this update, you were required to use the **openstack overcloud external-upgrade run --tags online_upgrade** command to perform online database updates when upgrading from RHOSP 15 to RHOSP 16.1. With this update, you can now use the **openstack overcloud external-update run --tags online_upgrade** command.

BZ#1889228

Before this update, cloned encrypted volumes were inaccessible when using the Block Storage (cinder) service with the Key Manager (barbican) service. With this update, cloned encrypted volumes are now accessible when using the Block Storage service with the Key Manager service.

BZ#1898484

Before this update, the connection data created by an iSCSI/LVM Block Storage back end was not stored persistently, which resulted in volumes not being accessible after a reboot. With this update, the connection data is stored persistently, and the volumes are accessible after a system reboot.

BZ#1899761

Before this update, when deployed at an edge site the Image (glance) service was not configured to access the Key Manager (barbican) service running on the central site's control plane. This resulted in the Image services running on edge sites being unable to access encryption keys stored in the Key Manager service.
With this update, Image services running on edge sites are now configured to access the encryption keys stored in the Key Manager service.

BZ#1901157

Before this update, in-place upgrades from Red Hat OpenStack Platform 13 to 16.1 in a TLS everywhere environment used an incorrect rabbitmq password for the novajoin container. This caused the novajoin container on the undercloud to function incorrectly, which caused any overcloud node that ran an upgrade to fail with the following error:

```
2020-11-24 20:01:31.569 7 ERROR join   File "/usr/lib/python3.6/site-
packages/amqp/connection.py", line 639, in _on_close
2020-11-24 20:01:31.569 7 ERROR join     (class_id, method_id), ConnectionError)
2020-11-24 20:01:31.569 7 ERROR join amqp.exceptions.AccessRefused: (0, 0): (403)
ACCESS_REFUSED - Login was refused using authentication mechanism AMQPLAIN. For detail
see the broker logfile.
```

With this update, the upgrade from RHOSP 13 to 16.1 uses the correct rabbitmq password in a TLS everywhere environment so that the framework for upgrades can complete successfully.

BZ#1902142

Before this update, when you configured the **collectd::plugin::virt::hostname_format** parameter with multiple values, director wrapped the values in double quotes. This caused the virt plugin to fail to load. With this update, when configuring **collectd::plugin::virt::hostname_format**, director no longer wraps multiple values in double quotes.

BZ#1906698

Before this update, live migration failed when upgrading a TLS everywhere environment with local ephemeral storage and **UseTLSTransportForNbd** set to "False". This occurred because the default value of the **UseTLSTransportForNbd** configuration had changed from "False" in RHOSP 13 to "True" in RHOSP 16.x, which resulted in the correct certifications not being included in the QEMU process containers.
With this update, director checks the configuration of the previously deployed environment for **global_config_settings** and uses it to ensure that the **UseTLSTransportForNbd** state stays the same in the upgrade as on previous deployment. If **global_config_settings** exists in the

configuration file, then director checks the configuration of the **use_tls_for_nbd** key. If **global_config_settings** does not exist, director evaluates the hieradata key **nova::compute::libvirt::qemu::nbd_tls**. Keeping the **UseTLSTransportForNbd** state the same in the upgraded deployment as on previous deployment ensures that live migration works.

BZ#1909795

Before this update, a rebase in python-network-runner from 0.1.7 to 0.2.2 in OSP 16.1.3 caused ML2 Networking using Ansible to no longer function.
With this update, python-networking-ansible is reverted to 0.1.7, and Ansible networking returns to a functioning state.

For more information, see https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.1/html/bare_metal_provisioning/ml2-networking-ansible.

BZ#1910854

Before this update, the Block Storage (cinder) NEC back end driver occasionally returned invalid data when initializing a volume connection, which could cause live migration to fail. With this update, the NEC driver has been fixed to reliably return valid connection data. Live migration no longer fails due to invalid volume connection data.

BZ#1921735

Before this update, the Block Storage (cinder) service would always assign newly created volumes with the default volume type, even when the volume was created from another source, such as an image, snapshot or another volume. This resulted in volumes created from another source having a different volume type from the volume type of the source.
With this update, the default volume type is assigned only after determining whether it should be assigned based on the volume type of the source. The volume type of volumes created from another source now match the volume type of the source.

BZ#1929275

Before this update, instances that were created on a RHOSP 13 environment with PowerFlex, VxFlex and ScaleIO volume attachments failed restarting after an upgrade to RHOSP 16.x. This was because the RHOSP 16.x Compute service uses a new PowerFlex driver connection property to access volume attachments, which is not present in the connection properties of volumes attached to instances running on a RHOSP 13 environment. With this update, the error is no longer thrown if this connection property is missing, and instances with PowerFlex volume attachments created on a RHOSP 13 environment continue to function correctly after upgrading to RHOSP 16.x.

3.5.3. Enhancements

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

BZ#1459187

Red Hat OpenStack Platform (RHOSP) 16.1 includes support 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. In RHOSP 16.1.2 this feature graduated from Technology Preview to full support.

BZ#1474394

Red Hat OpenStack Platform (RHOSP) 16.1 includes support for bare metal provisioning over an IPv6 provisioning network for BMaaS (Bare Metal as-a-Service) tenants. In RHOSP 16.1.2, this feature has graduated from Technology Preview to full support.

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#1640742

With this enhancement, you can configure NVDIMM Compute nodes to provide persistent memory for instances. Using this feature, you can make the PMEM available to instances as virtual PMEM (vPMEM) by creating and configuring PMEM namespaces on Compute nodes that have NVDIMM hardware. Cloud users can then create instances that request vPMEM when they need the instance content to be retained after it is shut down.
Note: Due to the availability of daxio package only for x86_64 architecture, this feature is supported only on x86_64 Compute nodes.

BZ#1793595

With this enhancement, you can deploy the Red Hat Ceph Storage (RHCS) Dashboard on edge sites in a distributed compute node (DCN) architecture.

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#1844615

This enhancement adds support for using Open Virtual Network (OVN) with Network Functions Virtualization infrastructure (NFVi) for new deployments. This includes support for the following features:

- OVN with OVS-DPDK
- OVN with SR-IOV
- OVN with OVS TC Flower offload

Note: Migration from ML2/OVS to ML2/OVN is not yet supported for NFV deployments.

BZ#1846019

This enhancement adds support for vlan transparency in the ML2/OVN mechanism driver with vlan and geneve network type drivers.
With vlan transparency, you can manage vlan tags by using instances on Networking (neutron) service networks. You can create vlan interfaces on an instance and use any vlan tag without affecting other networks. The Networking service is not aware of these vlan tags.

Notes:

- When using vlan transparency on a vlan type network, the inner and outer ethertype of the packets is 802.1Q (0x8100).
- The ML2/OVN mechanism driver does not support vlan transparency on flat provider networks.

BZ#1878191

With this enhancement, you can configure the format of the plugin instance for the collectd virt plugin by using the **ExtraConfig** parameter **collectd::plugin::virt::plugin_instance_format**. This allows more granular metadata to be exposed in the metrics label for virtual machine instances, such as on which host the instance is running.

BZ#1882058

This enhancement adds support for heterogeneous storage configurations at the edge. Operators can now deploy edge sites with storage and sites without storage within the same DCN deployment.

BZ#1891828

The Block Storage backup service sometimes needs access to files on the host that would otherwise not be available in the container running the service. This enhancement adds the CinderBackupOptVolumes parameter, which you can use to specify additional container volume mounts for the Block Storage backup service.

3.5.4. 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#1870199

The **virt-admin** tool is now available for you to use to capture logs for reporting RHOSP bugs. This tool is useful for troubleshooting all libvirt and QEMU problems, as the logs provide the communications between libvirt and QEMU on the Compute nodes. You can use **virt-admin** to set the libvirt and QEMU debug log filters dynamically, without having to restart the **nova_libvirt** container.
Perform the following steps to enable libvirt and QEMU log filters on a Compute node:

1. Log in to the **nova_libvirt** container on the Compute node:

```
$ sudo podman exec -it nova_libvirt /bin/bash
```

2. Specify the name and location of the log file to send **virt-admin** output to:

```
$ virt-admin daemon-log-outputs "1:file:/var/log/libvirt/libvirtd.log"
```

3. Configure the filters you want to collect logs for:

```
$ virt-admin daemon-log-filters \
 "1:libvirt 1:qemu 1:conf 1:security 3:event 3:json 3:file 3:object 1:util"
```

NOTE

When debugging issues with live migration, you must configure these filters on all source and destination Compute nodes.

4. Repeat your test. After debugging is complete, upload the **libvirtd.log** to a bug.

5. Disable the libvirt and QEMU log filters on the Compute nodes:

```
$ virt-admin daemon-log-filters ""
```

6. To confirm that the filters are removed, enter the following command:

```
$ virt-admin daemon-log-filters
```

This command returns an empty list when you have successfully removed the filters.

3.5.5. Known issues

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

BZ#1866479

There is currently a known issue with the mechanism that ensures the subscribed environments have the right DNF module stream set. The Advanced Virtualization repository is not always available in the subscription that the Ceph nodes use, which causes the upgrade or update of a Ceph node to fail when you try to enable virt:8.2. For more information on the known issue, see https://bugzilla.redhat.com/show_bug.cgi?id=1923887.
Workaround:

Override the **DnfStreams** parameter in the upgrade or update environment file to prevent the Ceph upgrade from failing:

```
parameter_defaults:
  ...
  DnfStreams: [{'module':'container-tools', 'stream':'2.0'}]
```

NOTE

The Advanced Virtualization DNF stream is not enforced when you use this workaround.

BZ#1925078

Systems that use UEFI boot and a UEFI bootloader in OSP13 might run into an UEFI issue that results in:

- **/etc/fstab** not being updated
- **grub-install** used incorrectly on EFI system

If your systems use UEFI, contact Red Hat Technical Support. For more information, see the Red Hat Knowledgebase solution FFU 13 to 16.1: Leapp fails to update the kernel on UEFI based systems and /etc/fstab does not contain the EFI partition.

BZ#1933268

There are currently known issues related to **[workarounds]/disable_native_luksv1** and **[workarounds]/rbd_volume_local_attach** configuration options. These options are provided only as a temporary workaround for known performance regressions within libgcrypt and librdb. These workaround options will be removed once the regressions are resolved in the underlying RHEL release used by RHOSP.
There are caveats associated with using either of these workaround options. Failure to adhere to these caveats can cause issues with RDB encrypted volumes. The caveats are as follows:

- A support exception must be granted from CEE before enabling these workaround options in your environment. A support exception is required to allow Red Hat to track use of the workaround and to help you remove these workaround options in the future when they are

disabled.

- You must enable the workaround options across all Compute nodes in a given environment or host aggregate.
- No move operations are supported between Compute nodes that have these workarounds enabled and those that do not have these workarounds enabled.
- All existing instances on a Compute node that you want to enable these workarounds on must be stopped or migrated off the node before you enable the workarounds. You can restart the instances when the Compute service has been restarted with the workarounds enabled.

BZ#1939419

If you use a Red Hat Ceph Storage subscription and have configured director to use the **overcloud-minimal** image for Red Hat Ceph Storage nodes, the upgrade of the operating system for Red Hat Ceph Storage nodes might fail due to a Leapp limitation.
To avoid this issue, after the **system_upgrade** run step, you must log in to the Red Hat Ceph Storage node to unset the RHEL minor release version, update to the latest available RHEL minor release version, and reboot the node.

If you use Red Hat Satellite Server to host RPM content for the Leapp upgrade, you must add the following 8.2 repositories to the Content View that you use:

- Red Hat Enterprise Linux 8 for x86_64 - AppStream (RPMs)

```
rhel-8-for-x86_64-appstream-rpms
x86_64 8.2
```

- Red Hat Enterprise Linux 8 for x86_64 - BaseOS (RPMs)

```
rhel-8-for-x86_64-baseos-rpms
x86_64 8.2
```

For more information, see https://bugzilla.redhat.com/show_bug.cgi?id=1936419

BZ#1942199

There is currently a known issue in **ceph-ansible** that prevents an ansible playbook from finishing successfully when it is used to deploy or update a Ceph Ganesha container that is configured to use an external, unmanaged Ceph cluster.
This issue causes deployments, minor updates, or major upgrades of overclouds that use Ceph Ganesha with the Shared File Systems service (manila), and that are configured to use an external Ceph cluster, to fail.

Workaround:

If you conduct an upgrade or minor update to the overcloud, use **ceph-ansible** version 4.0.49.1 or later if your environment is configured to use Ceph Ganesha with the Shared File Systems service and an external Ceph cluster. Under these conditions, do not attempt the upgrade with a version of **ceph-ansible** earlier than 4.0.49.1 installed on the undercloud.

3.6. RED HAT OPENSTACK PLATFORM 16.1.5 MAINTENANCE RELEASE - MARCH 31, 2021

3.6.1. Advisory list

This release includes the following advisories:

RHBA-2021:1052

Red Hat OpenStack Platform 16.1.5 bug fix and enhancement advisory

RHBA-2021:1053

Red Hat OpenStack Platform 16.1.5 containers bug fix advisory

RHBA-2021:1054

Red Hat OpenStack Platform 16.1.5 director images bug fix advisory

3.7. RED HAT OPENSTACK PLATFORM 16.1.6 MAINTENANCE RELEASE - MAY 27, 2021

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.

3.7.1. Advisory list

This release includes the following advisories:

RHBA-2021:2097::Red Hat OpenStack Platform 16.1.6 bug fix and enhancement advisory

RHSA-2021:2116::Moderate: Red Hat OpenStack Platform 16.1.6 (python-httplib2) security update

RHBA-2021:2117::Red Hat OpenStack Platform 16.1.6 containers bug fix advisory

RHBA-2021:2118::Red Hat OpenStack Platform 16.1.6 director images bug fix advisory

RHSA-2021:2119::Important: Red Hat OpenStack Platform 16.1.6 (tripleo-ansible) security update

3.7.2. Bug fix

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

BZ#1843788

This update fixes a bug that prevented **cinder list** from listing volumes when multiple filters were passed.

BZ#1868543

This update makes it possible to use OS::Heat:Delay resources in heat templates. Previously, a variable naming conflict caused an assertion error during attempted completion of an OS::Heat::Delay resource. A variable was renamed to eliminate the conflict.

BZ#1872314

When an instance is created, the Compute (nova) service sanitizes the instance display name to generate a valid host name when DNS integration is enabled in the Networking (neutron) service. Before this update, the sanitization did not replace periods ('.') in instance names, for example, 'rhel-

8.4'. This could result in display names being recognized as Fully Qualified Domain Names (FQDNs) which produced invalid host names. When instance names contained periods and DNS integration was enabled in the Networking service, the Networking service rejected the invalid host name, which resulted in a failure to create the instance and a HTTP 500 server error from the Compute service.

With this update, periods are now replaced by hyphens in instance names to prevent host names being parsed as FQDNs. You can continue to use free-form strings for instance display names.

BZ#1895045

This update fixes a bug that caused failure of validations before **openstack undercloud upgrade** in some cases. Before this upgrade, a lack of permissions needed to access the requested logging directory sometimes resulted in the following failures:

- Failure to log validation results
- Failure of the validation run
- Failure of artifacts collection from validation.
  This update adds a fallback logging directory. Validation results are logged and artifacts collected.

BZ#1905231

This update adds CHAP support to the Dell EMC PowerStore driver.

BZ#1910855

In prior releases, cinder NEC driver backups failed when the object was a snapshot. This occurred because the **snapshot** argument does not have the **volume_attachment** attribute. With this update, backups no longer refer to the **volume_attachment** attribute when the argument is **snapshot**.

BZ#1936419

This update fixes a configuration problem that caused Leapp upgrades to stop and fail while executing on a CephStorage node.
Previously, CephStorage nodes were incorrectly configured to consume OpenStack highavailability, advanced-virt, and fast-datapath repos during Leapp upgrades.

Now **UpgradeLeappCommand** options is configurable on a per-node basis, and uses the correct default for CephStorage nodes, and Leapp upgrades succeed for CephStorage nodes.

BZ#1939398

In prior releases, the SolidFire driver created a duplicate volume whenever it retried an API request. This led to unexpected behavior due to the accumulation of unused volumes.
With this update, the Block Storage service (cinder) checks for existing volume names before it creates a volume. When Block Storage service detects a read timeout, it immediately checks for volume creation to prevent invalid API calls. This update also adds the **sf_volume_create_timeout** option for the SolidFire driver so that you can set an appropriate timeout value for your environment.

BZ#1947474

This update fixes an issue that caused some API calls, such as create snapshot, to fail with an xNotPrimary error during workload re-balancing operations.
When SolidFire is under heavy load or being upgraded, the SolidFire cluster might re-balance cluster workload by automatically moving connections from primary to secondary nodes. Previously, some API calls failed with an xNotPrimary error during these workload balance operations and were not retried.

This update fixes the issue by adding the xNotPrimary exception to the SolidFire driver list of retryable exceptions.

3.7.3. 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#1546996

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

3.8. RED HAT OPENSTACK PLATFORM 16.1.7 MAINTENANCE RELEASE - DECEMBER 09, 2021

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.

3.8.1. Advisory list

This release includes the following advisories:

RHBA-2021:3762

Red Hat OpenStack Platform 16.1.7 bug fix and enhancement advisory

RHSA-2021:5070

Moderate: Red Hat OpenStack Platform 16.1 (python-django20) security update

RHSA-2021:5071

Moderate: Red Hat OpenStack Platform 16.1 (python-eventlet) security update

RHSA-2021:5072

Moderate: Red Hat OpenStack Platform 16.1 (etcd) security update

RHBA-2021:5073

Red Hat OpenStack Platform 16.1.7 containers bug fix advisory

RHBA-2021:5074

Red Hat OpenStack Platform 16.1.7 director images bug fix advisory

3.8.2. Bug fix

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

BZ#1906162

Before this update, the **appstream** and **baseos** repositories were always added to the repositories enabled by Red Hat Subscription Manager, with no way to override them. With this update, when you define the **$REG_REPOS** variable no base repositories are added. With this fix, you can fully control which repositories are added, but you must now include all repositories including the equivalent repository for **baseos** (and **appstream** when required).

BZ#1930255

In previous releases, in Red Hat OpenStack Platform (RHOSP) deployments that use the Dell EMC XtremIO driver, attach volume operations waited for a timeout if iSCSI or FC targets were not connected to a RHOSP host. This caused attach volume operations to fail.
This release adds port filtering support for the Dell EMC XtremIO driver to allow iSCSI or FC ports that are not in use to be ignored.

BZ#1938212

Before this update, the Shared File Systems service (manila) dashboard had dynamic form elements whose names could potentially cause the forms to become unresponsive. This meant that the creation of share groups, share networks, and shares within share networks did not function.
With this update, dynamic elements whose names might be problematic are encoded. Which means that creation of share groups, share networks, and shares within share networks functions normally.

BZ#1945306

In previous releases, if Dell EMC PowerStore ports were configured for multiple purposes, such as iSCSI, replication, incorrect REST filtering caused the cinder driver to report that no accessible iSCSI targets were found.
This release fixes the Dell EMC PowerStore REST filter functionality.

BZ#1947415

Before this update, a failure occurred when users wanted to delete the ***DEFAULT*** volume type.
With this update, you can delete the ***DEFAULT*** volume type when it is not set as the value of the **default_volume_type** parameter in the **cinder.conf** file. The default value of the **default_volume_type** parameter is ***DEFAULT*** so you must set it to an appropriate volume type, for example 'tripleo', so that you can delete the ***DEFAULT*** volume type.

BZ#1952574

Before this update, if your environment was deployed with a TLS-Everywhere architecture and it used the deprecated **authconfig** utility to configure authentication on your system, you had to configure your RHEL 8 system with the **authselect** utility. Without performing this action, the **leapp** process failed with the inhibitor named **Missing required answers in the answer file**. The workaround was to add **sudo leapp answer --section authselect_check.confirm=True --add** in the **LeappInitCommand** in the upgrades environment file. With this update, the configuration entry is no longer needed, and the upgrade now completes without intervention.

BZ#1959866

Before this update, during a tripleo validation on an OpenStack component, the following exception error occurred:

> Unhandled exception during validation run.

This error occurred because a variable in the code was referenced, but never assigned.

With this update, this problem has been fixed and validations run without this error.

BZ#1962365

Before this update, the Red Hat Enterprise Linux (RHEL) in-place upgrade tool, LEAPP, stalled because it encountered loaded kernel modules that are no longer provided in RHEL 8. Also, LEAPP upgraded RHEL to a version that is not supported by Red Hat OpenStack Platform (RHOSP). With this update, the manual configurations that you had to perform to workaround these two issues are no longer required. (For more information, see BZ1962365.

BZ#1974831

With this update, there is resolution to a problem that prevented the RHOSP Load-balancing service (octavia) to fail over load balancers with multiple failed amphorae.

BZ#1975790

Before this update, when a configuration change to a Load-balancing service amphora caused an haproxy reload, the process consumed a lot of memory that could lead to memory allocation errors. The issue was caused by the **lo** interface not being configured in the amphora-haproxy namespace in the amphora. With this update, the namespace issue has been corrected and the problem is resolved.

BZ#1977792

Before this update, there were unhandled exceptions during connection to iSCSI portals. For example, failures in **iscsiadm -m session**. This occurred because the **_connect_vol** threads can abort unexpectedly in some failure patterns, and this abort causes a hang in subsequent steps while waiting for results from **_connct_vol** threads.
With this update, any exceptions during connection to iSCSI portals are handled in the **_connect_vol** method correctly and avoids any unexpected abort without updating thread results.

BZ#1980829

Before this update, changes to **KernelArgs** parameters caused errors in the Red Hat OpenStack Platform (RHOSP) fast forward upgrade (FFU) process for version 13 to version 16:

- Duplicate entries appeared in **/etc/default/grub**.
- Duplicate entries appeared in the kernel command line.
- Nodes rebooted during the RHOSP upgrade.
  These errors were caused when the **KernelArgs** parameter, or the order of values in the string, changed or when a **KernelArgs** parameter was added.

  With this update, TripleO has added upgrade tasks in **kernel-boot-params-baremetal-ansible.yaml** to migrate from **TRIPLEO_HEAT_TEMPLATE_KERNEL_ARGS** to **GRUB_TRIPLEO_HEAT_TEMPLATE_KERNEL_ARGS**.

  This change was made to accommodate the Red Hat Enterprise Linux (RHEL) in-place upgrade tool, LEAPP, which is used to upgrade RHEL from version 7 to version 8, during the RHOSP version 13 to version 16 FFU process. LEAPP understands GRUB parameters only when the parameters start with **GRUB_** in **/etc/default/grub**.

  Despite this update, you must manually inspect each **KernelArgs** value to ensure that it matches the value for all hosts in the corresponding role.

  The **KernelArgs** value may come from the **PreNetworkConfig** implementation from either the default tripleo-heat-templates or third-party heat templates.

  If you find any mismatches, change the value of the **KernelArgs** parameter in the corresponding role to match the value of **KernelArgs** on the hosts. Perform these checks before running the **openstack overcloud upgrade prepare** command.

  You can use the following script to check **KernelArgs** values:

```
tripleo-ansible-inventory --static-yaml-inventory inventory.yaml
KernelArgs='< KernelArgs_ FROM_THT >'
ansible -i inventory.yaml ComputeSriov -m shell -b -a "cat /proc/cmdline | grep
'${KernelArgs}'"
```

BZ#1981652

Before this update, an optional feature of the RHOSP Load-balancing service (octavia), log offloading, was not correctly configured during deployment. As a result of this problem, the Load-balancing service was not receiving logs from the amphorae. This update resolves the issue.

BZ#1987104

Before this update, creating a volume from a snapshot of an encrypted volume could result in an unusable volume. When the destination volume is the same size as the source volume, creating an encrypted volume from a snapshot of an encrypted volume truncated the data in the new volume, which caused a size discrepancy.
With this update, the RBD back end accounts for the encryption header and does not truncate the data so that creating a volume from a snapshot of an encrypted volume does not cause the error.

BZ#1997351

Before this update, upgrading a Red Hat OpenStack Platform (RHOSP) 13 environment that has been deployed with ML2-OVN, to RHOSP 16.1 caused the upgrade process to fail on the Controller nodes due to an SELinux denial issue. With this update, the correct SELinux label is applied to OVN and resolves the issue. For more information, see the Red Hat Knowledgebase solution OVN fails to configure after reboot during OSP-13 → OSP-16.1 FFU.

BZ#2008976

Before this update, removal of the **python2** packages for the Red Hat Enterprise Linux (RHEL) in-place upgrade tool, LEAPP, was unsuccessful. This failure was caused by a DNF **exclude** option that retained the LEAPP packages. With this update, automation has now been included to ensure that the necessary LEAPP packages are successfully removed.

BZ#2015325

Before this update, an upgradable **mariadb-server** package in the RHEL repository caused the package manager to upgrade the **mariadb-server** package on the host, which interfered with the containerized **mariadb-server** that pre-exists on the same host. With this update, the Red Hat OpenStack Platform (RHOSP) director removes the **mariadb-server** package from any hosts that also have the containerized MariaDB, and the RHOSP FFU process continues.

3.8.3. Enhancements

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

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#1900500

The logic to detect the hypervisor hostname has been fixed and now returns the result consistent with **libvirt** driver in the Compute service (nova). With this fix, you no longer need to specify the **resource_provider_hypervisors** option when you use the guaranteed minimum bandwidth QoS feature.
With this update, a new option, **resource_provider_default_hypervisor**, has been added to the Modular Layer 2 with the Open Virtual Network mechanism driver (ML2/OVN) to replace the default hypervisor name. The option locates the root resource provider without giving a complete list of interfaces or bridges in the **resource_provider_hypervisors** option in case it has to be customized by the user. This new option is located in the **[ovs]** ini-section for the **ovs-agent**, and in the **[sriov_nic]** ini-section for the **sriov-agent**.

BZ#1930806

This enhancement adds the new **CinderRpcResponseTimeout** and **CinderApiWsgiTimeout** parameters to support tuning RPC and API WSGI timeouts in the Block Storage service (cinder). Default timeout values might not be adequate for large deployments and in situations where transactions might be delayed due to system load.
It is now possible to tune the RPC and API WSGI timeouts to prevent transactions prematurely timing out.

BZ#1956887

Previously, the **PluginInstanceFormat** parameter for collectd accepted only one of the following values: 'none', 'name', 'uuid', or 'metadata'. With this update, you can now specify more than one value for the **PluginInstanceFormat** parameter, resulting in more information being sent in the **plugin_instance** label of collectd metrics.

BZ#1959492

With this update, the **tripleo validator** command now accepts variables and environment variables in a key-value pair format. In past releases, only JSON dictionaries allowed environment variables.

```
openstack tripleo validator run \
[--extra-vars key1=<val1>[,key2=val2 --extra-vars key3=<val3>] \
| --extra-vars-file EXTRA_VARS_FILE] \
[--extra-env-vars key1=<val1>[,key2=val2 --extra-env-vars key3=<val3>]]
(--validation <validation_id>[,<validation_id>,...] | --group <group>[,<group>,...])
```

Example

```
$ openstack tripleo validator run --validation check-cpu,check-ram --extra-vars minimal_ram_gb=8
--extra-vars minimal_cpu_count=2
```

For the complete list of supported options, run:

```
$ openstack tripleo validator run --help
```

3.8.4. 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#1969895

With this update, the memory limit for the **collectd** container has been increased to 512 MB. When this limit is exceeded, the container restarts.

3.8.5. Known issues

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

BZ#1898198

Currently, there is a known issue where it is not possible to simulate certain real-life scenarios when the MAC-IP addresses of a port are unknown. The RHOSP Networking service (neutron) directly specifies the MAC-IP of a port even if DHCP or security groups are not configured.

Workaround: upgrade to RHOSP 16.1.7 and install ML2/OVN v21.03. If DHCP and port security are disabled, then the addresses field of a port does not include its MAC-IP address pairs, and ML2/OVN can use the MAC learning capabilities to send traffic only to the desired port.

3.9. RED HAT OPENSTACK PLATFORM 16.1.8 MAINTENANCE RELEASE - MARCH 23, 2022

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.

3.9.1. Advisory list

This release includes the following advisories:

RHSA-2022:0982

Important: Red Hat OpenStack Platform 16.1 (python-twisted) security update

RHSA-2022:0983

Moderate: Red Hat OpenStack Platform 16.1 (openstack-nova) security update

RHSA-2022:0984

Red Hat OpenStack Platform 16.1.8 director images bug fix advisory

RHSA-2022:0985

Red Hat OpenStack Platform 16.1.8 containers bug fix advisory

RHSA-2022:0987

Moderate: Red Hat OpenStack Platform 16.1 (numpy) security update

RHSA-2022:0986

Red Hat OpenStack Platform 16.1.8 bug fix and enhancement advisory

RHSA-2022:0988

Moderate: Red Hat OpenStack Platform 16.1 (golang-github-vbatts-tar-split) security update

RHSA-2022:0989

Moderate: Red Hat OpenStack Platform 16.1 (golang-qpid-apache) security update

RHSA-2022:0990

Moderate: Red Hat OpenStack Platform 16.1 (openstack-neutron) security update

3.9.2. Bug fix

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

BZ#1741453

Before this update, the OpenStack NFS driver blocked attempts to delete snapshots in an error state when snapshot support is disabled. New or existing snapshots are placed in an error state when snapshot support is disabled, but users could not remove these failed snapshots. With this update, users can now remove NFS snapshots in error status.

BZ#1815305

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

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

BZ#1910939

With this update, the telemetry healthchecks have been made more robust and the way the healthchecks are parsed has been simplified.
To get verbose mode when you run the healthcheck directly, run the command **sudo podman -u root -e "HEALTHCHECK_DEBUG=1" <container> /openstack/healthcheck**

BZ#1960639

Before this update, the OpenStack Storage service (cinder) GPFS SpectrumScale driver would not correctly detect whether the storage backend supported copy-on-write (COW) mode. As a result, the driver would disable COW features, such as the ability to rapidly create volumes from an image. This can cause some instances to time out when booting multiple instances simultaneously from an image.
With this update, the GPFS SpectrumScale driver properly detects COW support for storage backends.

BZ#1987957

Before this update, the PowerMax driver used a mechanism for storing and maintaining information on shared volume connections that did not work with previously created legacy volumes. This caused live migration to fail for volumes that were created before the PowerMax migration code was introduced. Now, the PowerMax live migration code is updated to work with legacy volumes so that live migrations do not fail.

BZ#1992159

Before this update, when creating a snapshot with PowerMaxOS 5978.711, REST experienced a payload response change and caused the device label to modify its format. The underlying data from the solutions enabler changed, and no longer contained a colon character (:). This, in turn, would cause an IndexError exception in the PowerMax Driver:

```
IndexError: list index out of range
```

With this update, the problem is resolved in PowerMaxOS 5978.711 and later.

BZ#1999634

This update fixes a bug that omitted details from the output of the **openstack volume backup list** command when the output exceeded 1000 lines.

BZ#1999901

This update fixes a unicode escaping bug that caused Horizon to crash when passwords were changed while the Horizon language was set to Japanese. Changing the password no longer causes Horizon to crash.

BZ#2023413

Before this update, os-brick did not include a **[global]** section to contain the options it sets in a temporary configuration file, which is a requirement with Octopus (release 15.2.0+). As a result, connection information could not be found when using os-brick and a Ceph Octopus or later client, and a connection to the Ceph storage backend could not be established. Now, the connection options are included under a '[global]' section in the temporary configuration file. This fix is backward compatible to the Hammer release (0.94.0+) of Ceph.

BZ#2029608

This update corrects an error that prevented the proper use of the Cinder **powermax_port_groups** parameter.

3.9.3. Enhancements

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

BZ#1939964

Enable the experimental **rsyslog reopenOnTruncate** to ensure that rsyslog immediately recognizes when a logrotation happens on a file. The setting affects every service configured to work with rsyslog.
With **rsyslog reopenOnTruncate** disabled, rsyslog waits for a log file to fill to its original capacity before consuming any additional logs.

BZ#1949168

This enhancement prepares your environment for update of the metrics_qdr service to a newer AMQ Interconnect release, which requires import of the CA certificate contents from the Service Telemetry Framework (STF) deployment. Changes are not yet required by administrators when deploying or updating Red Hat OpenStack Service Platform (RHOSP) as the metrics_qdr service has not yet been updated. This functionality is in preparation of the metrics_qdr service update in a future release.
The following procedure will be required once https://bugzilla.redhat.com/show_bug.cgi?id=1949169 has shipped.

This update corrects this problem by providing a new Orchestration service (heat) parameter, **MetricsQdrSSLProfiles**.

To obtain a Red Hat OpenShift TLS certificate, run the following commands:

```
$ oc get secrets
$ oc get secret/default-interconnect-selfsigned -o jsonpath='{.data.ca\.crt}' | base64 -d
```

Add the **MetricsQdrSSLProfiles** parameter with the contents of your Red Hat OpenShift TLS certificate to a custom environment file:

```
MetricsQdrSSLProfiles:
  - name: sslProfile
    caCertFileContent: |
      -----BEGIN CERTIFICATE-----
      ...
      TOpbgNlPcz0sIoNK3Be0jUcYHVMPKGMR2kk=
      -----END CERTIFICATE-----
```

Then, redeploy your overcloud with the **openstack overcloud deploy** command.

BZ#1969999

With this update you can set QoS maximum bandwidth limit, egress direction rules on hardware-offloaded ports in a ML2/OVS deployment. Set the policy using the normal QoS policy/rules methods.
Note that the backend uses **ip link** commands to enforce the policy instead of the normal OVS QoS engine, because the OVS **meter** action cannot be offloaded. See meter action is not offloaded .

BZ#1984873

With this update, the **LeapActorsToRemove** heat parameter is introduced so that you can remove specific actors from the leapp process if those actors inhibit the upgrade. The **LeapActorsToRemove** heat parameter is role-specific for flexibility.

BZ#1992622

This feature enables Red Hat OpenStack Platform 16.1 to consume an external Red Hat Ceph Storage version 5 cluster.

BZ#2052411

As of this release, the Red Hat supported method of updating OVN is aligned to the upstream OVN updgrade steps.

3.9.4. 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#1984095

With this update the **CollectdContainerAdditionalCapAdd** variable is added to the deployment tool. This variable is a comma separated list of additional collectd container capabilities.

3.9.5. Known issues

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

BZ#2132151

A known issue causes Neutron to fail to start after an update to RHOSP 16.1.8. If you then get Neutron to start, the OVN databases are unstable.
A fix is planned for RHOSP 16.1.9. Red Hat recommends that you wait to update directly to 16.1.9 if possible.

A hot fix is available for 16.1.8. If you have an urgent need to update to RHOSP 16.1.8, contact Red Hat Global Support Services to see if your environment is compatible with the hot fix.

To track this issue, see https://bugzilla.redhat.com/show_bug.cgi?id=2125824.

3.9.6. Removed functionality

BZ#1996865

Before this release, the collectd container would fail to start on compute nodes because a dpdk-telemetry collectd configuration file was being automatically created despite there being no dpdk-telemetry plugin installed.
As of this release, dpdk_telemetry configuration files have been removed from the the collectd container.

3.10. RED HAT OPENSTACK PLATFORM 16.1.9 MAINTENANCE RELEASE - DECEMBER 7, 2022

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.

3.10.1. Advisory list

This release includes the following advisories:

RHEA-2022:8858

Red Hat OpenStack Platform 16.1.9 director images

RHBA-2022:8795

Red Hat OpenStack Platform 16.1.9 bug fix and enhancement advisory

RHEA-2022:8859

Red Hat OpenStack Platform 16.1.9 director image RPMs

RHSA-2022:8860

Moderate: Red Hat OpenStack Platform 16.1.9 (protobuf) security update

RHSA-2022:8861

Moderate: Red Hat OpenStack Platform 16.1.9 (numpy) security update

RHSA-2022:8862

Moderate: Red Hat OpenStack Platform 16.1.9 (puppet) security update

RHSA-2022:8863

Moderate: Red Hat OpenStack Platform 16.1.9 (python-paramiko) security update

RHSA-2022:8864

Moderate: Red Hat OpenStack Platform 16.1.9 (python-ujson) security update

RHSA-2022:8865

Moderate: Red Hat OpenStack Platform 16.1.9 (python-XStatic-Bootstrap-SCSS) security update

RHSA-2022:8866

Moderate: Red Hat OpenStack Platform 16.1.9 (python-XStatic-Angular) security update

RHSA-2022:8867

Low: Red Hat OpenStack Platform 16.1.9 (rabbitmq-server) security update

RHSA-2022:8868

Moderate: Red Hat OpenStack Platform 16.1.9 (python-scciclient) security update

RHSA-2022:8869

Moderate: Red Hat OpenStack Platform 16.1.9 (puppet-firewall) security update

RHSA-2022:8870

Moderate: Red Hat OpenStack Platform 16.1.9 (openstack-neutron) security update

RHBA-2022:8871

Updated Red Hat OpenStack Platform 16.1.9 container images

RHSA-2022:8872

Important: Red Hat OpenStack Platform 16.1.9 (python-django20) security update

RHSA-2022:8873

Moderate: Red Hat OpenStack Platform 16.1.9 (python-oslo-utils) security update

RHSA-2022:8796

Moderate: Red Hat OpenStack Platform 16.1.9 (openstack-tripleo-heat-templates) security update

RHSA-2022:8874

Moderate: Red Hat OpenStack Platform 16.1.9 (openstack-barbican) security update

3.10.2. Bug Fix

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

BZ#1802263

Before this update, if you imported a backup record for a backup ID that currently existed, the import operation would correctly fail, but the existing backup record would incorrectly be deleted. With this update, the existing backup record is not deleted under this scenario.

BZ#1951485

Before this update, NetApp ONTAP Block Storage (cinder) driver QoS policy groups were deleted when the associated volume was moved. With this update, QoS policy groups are associated permanently to the LUN or file that represents the volume.

BZ#1961162

Before this update, a nonexistent gateway address was configured on the load-balancing management network. This caused excessive Address Resolution Protocol (ARP) requests on the load-balancing management network.

BZ#1968228

Before this update, the API that the Shared File Systems service (manila) uses to provision storage on NetApp ONTAP All Flash Fabric-Attached (AFF) storage systems caused Shared File Systems service shares to be thinly provisioned. The API did not enforce space guarantees, even when requested through the Shared File Systems service share type. With this update, the driver sets appropriate parameters for the NetApp ONTAP 9 API to work with AFF storage as well as traditional FAS storage systems. The API enforces space guarantees on NetApp ONTAP storage through the Shared File Systems service share types.

BZ#1977322

Before this update, a race condition occurred when the Compute service (nova) requested the Block Storage service (cinder) to detach a volume and there was an external request to delete the volume. The race condition resulted in the volume failing to detach, the volume being deleted, and the Compute service being unable to remove the non-existent volume. With this update, the race condition is resolved.

BZ#1996088

Before this update, python-octaviaclient did not display the full list of load balancers when the user had more than 1,000 load balancers. With this update, the OpenStack Load-balancing service (Octavia) displays all load balancers.

BZ#1996756

Before this update, members in the ERROR operating status might have been updated briefly to ONLINE during a Load Balancer configuration change. With this update, the issue is fixed.

BZ#2026029

Before this update, the secret:delete policy in the Key Manager service (barbican) only allowed users with the Creator role to delete a secret if they were the same user that created the secret. This limitation impacted encrypted workflows because of the mismatch in policy, for example, the Block Storage service (cinder) allows users with a role assignment on the project to delete encrypted volumes. However, the Key Manager service responded with an authorization error because not all users were allowed to delete the secret. With this update, the secret:delete policy in the Key Manager service has been changed to allow users with the Creator role to delete any secret that belongs to the project, not just the ones that they created. All users allowed to delete a Block Storage service encrypted volume are also allowed to delete the associated secret.

BZ#2027544

Before this update, if there were repeated transient connectivity issues between the ironic-conductor service and a remote Baseboard Management Controller (BMC) using the Redfish

hardware type when session authentication was used, the intermittent loss of connectivity could collide with a point where authentication was retried due to the in-memory credentials expiring. If this collision occurred, there was a loss of overall connectivity, which persisted due to the internal session cache built into the openstack-ironic-conductor service. With this update, support to detect and renegotiate in cases of this error were added to the Python DMTF Redfish library, sushy, and the openstack-ironic service. Intermittent connectivity failures colliding with session credential re-authentication no longer results in a complete loss of ability to communicate with the BMC until the openstack-ironic-conductor service is restarted.

BZ#2033953

Before this update, the machine-config-operator passed an afterburn systemd unit to new machines that set the hostname based on the user data passed through the Compute service (nova) metadata service. In some cases, for example, bare metal, the instance did not have connectivity to Compute service metadata. With this update, the afterburn systemd unit attempts to fetch data from configdrive first and then falls back to the Compute service metadata service. The hostname of instances is set, irrespective of the reachability of the Compute service metadata service.

BZ#2034095

Before this update, NTP validation did not occur during deployments. Some users reported issues with cloud authentication failing with invalid tokens due to time not being synchronized between nodes. With this update, NTP synchronization validation during deployment has been re-enabled. Hosts must be able to connect to the defined NTP server list. If you previously performed a deployment with invalid or unreachable NTP servers, after update, the deployment might fail when NTP is validated. Ensure that you have valid and reachable NTP servers before updating.

BZ#2040697

Before this update, the provisioning status of a load balancer was set to ERROR too early when an error occurred, making the load balancer mutable before the execution of the tasks for these resources was finished. With this update, the issue is fixed.

BZ#2057604

Before this update, the Load-balancing services (octavia) were restarted many times during deployments or updates. With this update, the services are restarted only when required, preventing potential interruptions of the control plane.

BZ#2063031

Before this update, systemd stopped the Load-balancing services (octavia) during shutdown, leaving resources in the PENDING_UPDATE status. With this update, the graceful shutdown duration of the Load-balancing services is increased, preventing the services from being stopped by systemd.

BZ#2064709

When a load balancer is created in a tenant network with a Virtual IP (VIP) and members, and the tenant network is connected to a router that is connected to the provider network, the Open Virtual Network (OVN) load balancer is associated with the OVN logical router. If the 'router' option was used for nat-addresses, ovn-controller sent GARP packets for that VIP on the provider network. As there was nothing to prevent different tenants in OpenStack from creating a subnet with the same Classless Inter-Domain Routing (CIDR) number and a load balancer with the same VIP, there could be several ovn-controllers generating GARP packets on the provider network for the same IP, each one with the MAC of the logical router port belonging to each tenant. This setup could be an issue for the physical network infrastructure. With this update, a new option (exclude-lb-vips-from-garp) is added in OVN[1] on the router gateway port. This flag ensures that no GARP packets are sent for the load balancer VIPs.

BZ#2078377

Before this update, the Virtual IP (VIP) address of UDP-only load balancers in active-standby mode was not reachable. With this update, the issue is fixed.

BZ#2089382

Before this update, block device mapping updates by the libvirt driver on the destination host were not persisted during live migration. With specific storage back ends or configurations, for example, when using the **n[workarounds]/rbd_volume_local_attach=True** config option, certain operations on volume attachments, for example detaching, after a live migration did not work. With this update, the Compute service (nova) correctly persists any block device mapping updates done by the libvirt driver on the destination host. Operations on affected volumes, such as detaching, succeed after live migration.

BZ#2096387

Before this update, a SELinux issue triggered errors when using the ICMP monitor in the Load-balancing service (octavia) amphora driver. With this update, the SELinux issue is fixed.

BZ#2100879

Before this update, dogpile.cache support for dead_retry and socket_timeout was not implemented for the memcached back end. The oslo.cache mechanism filled the arguments dictionary with values for dead_retry and socket_timeout, but dogpile.cache ignored the values so the defaults of 30s for dead_retry and 3s for socket_timeout were used. When using dogpile.cache.memcached as the cache back end on the Identity service (keystone), and then taking down one of the memcached instances, the memcache server objects set their deaduntil value to 30 seconds in the future. When a request came in to an API server with two memcached servers configured, one of which was unroutable, it took approximately 15 seconds for it to try each of those servers in each thread it created and reach the three-second socket timeout limit every time it encountered the one that was down. By the time the user issued another request, the deaduntil value was reached and the whole cycle was repeated. With this update, dogpile.cache consumes dead_retry and socket_timeout arguments passed by oslo.cache.

BZ#2103971

Before this update, the ceilometer-agent-compute container could not read the /var/run/libvirt directory because of an improper volume mount to /var/run/libvirt in the ceilometer-agent-compute container, resulting in the inability to poll for CPU metrics on Compute nodes. With this update, the appropriate global permissions have been applied to the /var/run/libvirt directory, and you can poll for CPU telemetry with the ceilometer-agent-compute container on the Compute nodes. CPU telemetry data is available through the Compute service (nova).

BZ#2122925

Before this update, it was possible to add members without stating which subnet they belonged to, but they should be in the same subnet as the Virtual IP (VIP) port. If the subnet of the members is different to the VIP subnet, the members are created but incorrectly configured because there is no connectivity to them. With this update, members without a subnet are only accepted if the IP of the member belongs to the Classless Inter-Domain Routing (CIDR) number of the VIP subnet, as that is the subnet associated to the load balancer used to obtain the subnet for the members that do not have it. Member creation without a subnet is rejected if its IP does not belong to the VIP subnet CIDR.

BZ#2123225

Before this update, Conntrack was enabled in the Amphora VM for any type of packet, but it is only required for the User Datagram Protocol (UDP) and Stream Control Transmission Protocol (SCTP). With this update, Conntrack is now disabled for Transmission Control Protocol (TCP) flows, preventing some performance issues when a user generates a lot of connections that fill the Conntrack table.

BZ#2125824

This update in RHOSP 16.1.9 fixes a bug that causes the Networking service (neutron) to fail to start after an update to RHOSP 16.1.8 and also causes OVN database instability after updates to RHOSP 16.1.8.
Instead of updating to RHOSP 16.1.8, update directly to RHOSP 16.1.9.

BZ#2129310

Before this update, a do_sync_check operation could result in the incorrect deletion of non-temporary snapshots from a volume because there was no check for non-temporary snapshots deletion during the do_sync_check operation. With this update, there is a check to determine if a snapshot must be deleted. The do_sync_check operation does not perform unnecessary non-temporary snapshot deletions.
Before this update, there was a case mismatch in the conditional while checking if a storage group was a child of a parent storage group. While modifying the storage group, errors indicated that the parent storage group already contained the child storage group. With this update, the patterns used in the conditional are not case-sensitive and you can modify the storage group successfully.

BZ#2130078

Before this update, the libvirt service started after the ceilometer-agent-compute service and the ceilometer-agent-compute service did not communicate with libvirt, resulting in missing libvirt metrics. With this update, the ceilometer-agent-compute service starts after the libvirt service and can poll libvirt metrics without "Permission denied" errors.

BZ#2130849

Before this update, a Telemetry service (ceilometer) user had insufficient privileges to poll objects from the Object Storage service (swift). The Object Storage service client did not allow the Telemetry service user to fetch object details. With this update, the Telemetry service user is associated with the ResellerAdmin role.
Execute the following command to workaround this issue manually:

```
$ openstack role add --user ceilometer --project service ResellerAdmin
```

The associated Telemetry service user can poll Object Storage service object metrics successfully.

BZ#2138184

RHSA-2022:6969 introduced the process to clean up files in the /var/lib/mistral directory in the undercloud but the process consistently failed when the Load-balancing service (octavia) or Red Hat Ceph Storage was enabled because these services created additional directories, which the cleanup process could not properly remove. Some deployment actions, such as scale out, consistently failed if the Load-balancing service or Ceph Storage was enabled. With this update, Mistral no longer executes the cleanup. Users must manually delete files if they want to enforce the reduced permission of the files in the /var/lib/mistral directory. Deployment actions no longer fail because of a permission error.

3.10.3. Enhancements

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

BZ#1917356

With this update, director supports specifying overrides for NVSv4 ID mapping when using a CephFS-NFS back end with the Shared File Systems service (manila). Ceph-NFS with the Shared File Systems service only allows client access through NFSv4.1+. With NFSv4.1, usernames and group names are sent over the wire and translated by both the server and the client. Deployers might want to customize their domain settings to better represent organization users who can access Shared File Systems service shares from multiple clients. Director supports customizing NFS ID mapping settings through these parameters:

- ManilaCephFSNFSIdmapOverrides: Allows specifying configuration objects for override with the default idmapd.conf file used by the NFS service

- ManilaCephFSNFSIdmapConf: Allows specifying a custom idmapd.conf file for the NFS service

BZ#1945334

With this update, the Rsyslog environment configuration supports an array of Elasticsearch targets. In previous releases, you could only specify a single target. You can now specify multiple Elasticsearch targets as a list of endpoints to send logs.

BZ#1982268

With this update, the **port_security** parameter of the Load-balancing service (octavia) management network is now enabled.

BZ#2022040

With this update, you can now migrate an ML2/OVS deployment with the iptables_hybrid firewall driver to ML2/OVN.

BZ#2070629

With this update, if the uplink_status_propagation extension is enabled, all single root I/O virtualization (SR-IOV) ports created before the extension enablement set the virtual function (VF) link state to 'auto'. Before this update, an SR-IOV port set the link state to enabled or disabled.

3.10.4. Known Issues

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

BZ#1574431

Currently, quota commands do not work as expected in the Block Storage service (cinder). The Block Storage CLI does not verify if the project ID that you specify is valid. Therefore, you can use the Block Storage CLI to create quota entries with invalid project IDs. These quota entries 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#2001012

You can use Role-based Access Control (RBAC) to share security groups between projects. However, you cannot use the **--security-group** argument to assign an RBAC-shared security group when you launch an instance. If you try to assign the RBAC-shared security group by using the **--security-group argument** in the **openstack server create** command, the Compute service (nova) does not find the security group and fails to create the instance. This is because the Compute service does not check for security groups that are shared through RBAC. It only checks if the security group specified with the **--security-group** argument exists in the project that you created the instance in.
Workaround: Create a port and assign the security group to the port. Use the **--nic** argument in the **openstack server create** command to specify the port. The Compute service does not try to create the port in the Networking service (neutron), therefore it does not check security groups.

For example:

```
$ openstack port create --network net1 \
  --security-group \
  5ba835b7-22b0-4be6-bdbe-e0722d1b5f24 shared-sg-port

$ openstack server create \
  --image cirros-0.5.1-x86_64-disk  \
  --flavor m1.tiny \
  --port shared-sg-port vm-with-shared-sg
```

BZ#2076884

There is currently a known issue when live migrating instances with CPUs that are incompatible with the destination host CPUs. As a workaround, you can skip the Compute service CPU comparison check on the destination host before migrating an instance, because libvirt (QEMU >= 2.9 and libvirt >= 4.4.0) correctly handles the CPU compatibility checks on the destination host during live migration.
Workaround: Before performing instance live migration, add the following configuration in the **nova.conf** file of each affected Compute node:

```
[workarounds]
skip_cpu_compare_on_dest = True
```

3.10.5. Removed Functionality

BZ#2101949

In Red Hat OpenStack Platform (RHOSP) 16.1.9, the collectd processes plugin is removed from the default list of plugins. Loading the plugin can cause flooding issues and does not provide value when running in a containerized environment because it only recognizes the collectd and sensubility processes rather than the expected system processes. Bug fixes and support will be provided through the end of the 16.1.9 lifecycle but no new feature enhancements will be made.

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_rsync** container runs in unprivileged mode. This makes the **swift_rsync** 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 enhanceded 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-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)

4.2. RHBA-2020:3542 – RED HAT OPENSTACK PLATFORM 16.1.1 GENERAL AVAILABILITY ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2020:3542. Further information about this advisory is available at link: https://access.redhat.com/errata/RHBA-2020:3542.html.

Changes to the openstack-tripleo component:

- The overcloud deployment steps included an older Ansible syntax that tagged the **tripleo-bootstrap** and **tripleo-ssh-known-hosts** roles as **common_roles**. This older syntax caused Ansible to run tasks tagged with the **common_roles** when Ansible did not use the **common_roles** tag. This syntax resulted in errors during the 13 to 16.1 **system_upgrade** process.
  This update uses a newer syntax to tag the **tripleo-bootstrap** and **tripleo-ssh-known-hosts** roles as **common_roles**. Errors do not appear during the 13 to 16.1 **system_upgrade** process and you no longer include the **--playbook upgrade_steps_playbook.yaml** option to the **system_upgrade** process as a workaround. (BZ#1851914)

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

- This update fixes a GRUB parameter naming convention that led to unpredictable behaviors on compute nodes during leapp upgrades.
  Previously, the presence of the obsolete "TRIPELO" prefix on GRUB parameters caused problems.

  The file /etc/default/grub has been updated with GRUB for the tripleo kernel args parameter so that leapp can upgrade it correctly. This is done by adding "upgrade_tasks" to the service "OS::TripleO::Services::BootParams", which is a new service added to all roles in the roles_data.yaml file. (BZ#1858673)

- This update fixes a problem that caused baremetal nodes to become non-responsive during Leapp upgrades.
  Previously, Leapp did not process transient interfaces like SR-IOV virtual functions (VF) during migration. As a result, Leapp did not find the VF interfaces during the upgrade, and nodes entered an unrecoverable state.

  Now the service "OS::TripleO::Services::NeutronSriovAgent" sets the physical function (PF) to remove all VFs, and migrates workloads before the upgrade. After the successful Leapp upgrade, os-net-config runs again with the "--no-activate" flag to re-establish the VFs. (BZ#1866372)

- This director enhancement automatically installs the Leapp utility on overcloud nodes to prepare for OpenStack upgrades. This enhancement includes two new Heat parameters: LeappRepoInitCommand and LeappInitCommand. In addition, if you have the following repository defaults, you do not need to pass UpgradeLeappCommandOptions values. --enablerepo rhel-8-for-x86_64-baseos-eus-rpms --enablerepo rhel-8-for-x86_64-appstream-eus-rpms --enablerepo rhel-8-for-x86_64-highavailability-eus-rpms --enablerepo advanced-virt-for-rhel-8-x86_64-rpms --enablerepo ansible-2.9-for-rhel-8-x86_64-rpms --enablerepo fast-datapath-for-rhel-8-x86_64-rpms

  (BZ#1845726)

- If you do not set the **UpgradeLevelNovaCompute** parameter to **''**, live migrations are not possible when you upgrade from RHOSP 13 to RHOSP 16. (BZ#1849235)

- This update fixes a bug that prevented the successful deployment of transport layer security (TLS) everywhere with public TLS certifications. (BZ#1852620)

- 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#1853275)

- 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. (BZ#1853433)

Changes to the openstack-tripleo-validations component:

- This update fixes a Red Hat Ceph Storage (RHCS) version compatibility issue that caused failures during upgrades from Red Hat OpenStack platform 13 to 16.1. Before this fix, validations performed during the upgrade worked with RHCS3 clusters but not RHCS4 clusters. Now the validation works with both RHCS3 and RHCS4 clusters. (BZ#1852868)

Changes to the puppet-tripleo component:

- 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#1850991)

4.3. RHSA-2020:4283 – RED HAT OPENSTACK PLATFORM 16.1.2 GENERAL AVAILABILITY ADVISORY

The bugs contained in this section are addressed by advisory RHSA-2020:4283. Further information about this advisory is available at link: https://access.redhat.com/errata/RHSA-2020:4283.html.

Bug Fix(es):

- This update includes the following bug fix patches related to fully qualified domain names (FQDN).

  - *Kaminario Fix unique_fqdn_network option*
    Previously, the Kaminario driver accepted the unique_fqdn_network configuration option in the specific driver section. When this option was moved, a regression was introduced: the parameter was now only used if it was defined in the shared configuration group.

    This patch fixes the regression and makes it possible to define the option in the shared configuration group as well as the driver specific section.

  - *HPE 3PAR Support duplicated FQDN in network*

The 3PAR driver uses the FQDN of the node that is doing the attach as an unique identifier to map the volume.

Because the FQDN is not always unique, in some environments the same FQDN can be found in different systems. In those cases, if both try to attach volumes, the second system will fail.

For example, this could happen in a QA environment where VMs share names like controller-.localdomain and compute-0.localdomain.

This patch adds the **unique_fqdn_network** configuration option to the 3PAR driver to prevent failures caused by name duplication between systems. (BZ#1721361) (BZ#1721361)

- This update makes it possible to run the Brocade FCZM driver in RHOSP 16.
The Brocade FCZM vendor chose not to update the driver for Python 3, and discontinued support of the driver past the Train release of OpenStack [1]. Red Hat OpenStack (RHOSP) 16 uses Python 3.6.

  The upstream Cinder community assumed the maintenance of the Brocade FCZM driver on a best-effort basis, and the bugs that prevented the Brocade FCZM from running in a Python 3 environment (and hence in RHOSP 16) have been fixed.

  [1] https://docs.broadcom.com/doc/12397527 (BZ#1848420)

- This update fixes a problem that caused volume attachments to fail on a VxFlexOS cinder backend.
Previously, attempts to attach a volume on a VxFlexOS cinder backend failed because the cinder driver for the VxFlexOS back end did not include all of the information required to connect to the volume.

  The VxFlexOS cinder driver has been updated to include all the information required in order to connect to a volume. The attachments now work correctly. (BZ#1862213)

- This enhancement introduces support for the revert-to-snapshot feature with the Block Storage (cinder) RBD driver. (BZ#1702234)

- 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#1808583)
- 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_volid** 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, 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#1826741)

- Before this update, when you created a volume from a snapshot, the operation could fail because the Block Storage service (cinder) would try to assign the default volume type to the new volume instead of inferring the correct volume type from the snapshot. With this update, you no longer have to specify the volume type when you create a volume. (BZ#1843789)

- This enhancement adds a new driver for the Dell EMC PowerStore to support Block Storage service back end servers. The new driver supports the FC and iSCSI protocols, and includes these features:
  - Volume create and delete
  - Volume attach and detach
  - Snapshot create and delete
  - Create volume from snapshot
  - Get statistics on volumes
  - Copy images to volumes
  - Copy volumes to images
  - Clone volumes
  - Extend volumes
  - Revert volumes to snapshots (BZ#1862541)

4.4. RHEA-2020:4284 – RED HAT OPENSTACK PLATFORM 16.1.2 GENERAL AVAILABILITY ADVISORY

The bugs contained in this section are addressed by advisory RHEA-2020:4284. Further information about this advisory is available at link: https://access.redhat.com/errata/RHEA-2020:4284.html.

Changes to the openstack-nova component:

- This bug fix enables you to boot an instance from an encrypted volume when that volume was created from an image that in turn was created by uploading an encrypted volume to the Image Service as an image. (BZ#1879190)

Changes to the openstack-octavia component:

- 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 <heath_monitor_id>**

  For more information, search for "loadbalancer healthmonitor" in the Command Line Interface Reference. (BZ#1837316)

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

- A known issue causes the migration of Ceph OSDs from Filestore to Bluestore to fail. In use cases where the **osd_objectstore** parameter was not set explicitly when you deployed OSP13 with RHCS3, the migration exits without converting any OSDs and falsely reports that the OSDs are already using Bluestore. For more information about the known issue, see https://bugzilla.redhat.com/show_bug.cgi?id=1875777
  As a workaround, perform the following steps:

  1. Include the following content in an environment file:

     ```
     parameter_defaults:
       CephAnsibleExtraConfig:
         osd_objectstore: filestore
     ```

  2. Perform a stack update with the **overcloud deploy --stack-only** command, and include the new or existing environment file that contains the **osd_objectstore** parameter. In the following example, this environment file is **<osd_objectstore_environment_file>**. Also include any other environment files that you included during the converge step of the upgrade:

     ```
     $ openstack overcloud deploy --stack-only \
       -e <osd_objectstore_environment_file> \
       -e <converge_step_environment_files>
     ```

  3. Proceed with the FileStore to BlueStore migration by using the existing documentation. See https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.1/html/framework_for_upgrades_13_to_16.1/OSD-migration-from-filestore-to-bluestore
     Result: The Filestore to Bluestore playbook triggers the conversion process, and removes and re-creates the OSDs successfully. (BZ#1733577)

- Inadequate timeout values can cause an overcloud deployment to fail after four hours. To prevent these timeout failures, set the following undercloud and overcloud timeout parameters:

- Undercloud timeouts (seconds):

  Example

  ```
  parameter_defaults:
    TokenExpiration: 86400
    ZaqarWsTimeout: 86400
  ```

- Overcloud deploy timeouts (minutes):

  Example

  ```
  $ openstack overcloud deploy --timeout 1440
  ```

  The timeouts are now set. (BZ#1792500)

- Currently, you cannot scale down or delete compute nodes if Red Hat OpenStack Platform is deployed with TLS-e using tripleo-ipa. This is because the cleanup role, traditionally delegated to the undercloud as localhost, is now being invoked from the mistral container.
  For more information, see https://access.redhat.com/solutions/5336241 (BZ#1866562)

- This update fixes a bug that prevented the distributed compute nodes (DCN) compute servcie from accessing the glance service.
  Previously, distributed compute nodes were configured with a glance endpoint URI that specified an IP address, even when deployed with internal transport layer security (TLS). Because TLS requires the endpoint URI to specify a fully qualified domain name (FQDN), the compute service could not access the glance service.

  Now, when deployed with internal TLS, DCN services are configured with glance endpoint URI that specifies a FQDN, and the DCN compute service can access the glance service. (BZ#1873329)

- This update introduces support of Distributed Compute Nodes TLS everywhere with Triple IPA. (BZ#1874847)

- The update introduces support of Neutron routed provider networks with RH-OSP Distributed Compute Nodes (BZ#1874863)

- This update adds support for encrypted volumes and images on distributed compute nodes (DCN).
  DCN nodes can now access the Key Manager service (barbican) running in the central control plane.

  

  NOTE

  This feature adds a new Key Manager client service to all DCN roles. To implement the feature, regenerate the **roles.yaml** file used for the DCN site's deployment.

  For example:

  ```
  $ openstack overcloud roles generate DistributedComputeHCI
  DistributedComputeHCIScaleOut -o ~/dcn0/roles_data.yaml
  ```

  Use the appropriate path to the roles data file. (BZ#1852851)

- Before this update, to successfully run a leapp upgrade during the fast forward upgrade (FFU) from RHOSP 13 to RHOSP16.1, the node where the Red Hat Enterprise Linux upgrade was occurring had to have the **PermitRootLogin** field defined in the ssh config file (**/etc/ssh/sshd_config**).
  With this update, the Orchestration service (heat) no longer requires you to modify **/etc/ssh/sshd_config** with the **PermitRootLogin** field. (BZ#1855751)

- This enhancement adds a new driver for the Dell EMC PowerStore to support Block Storage service back end servers. (BZ#1862547)

Changes to the openstack-tripleo-validations component:

- This update safeguards against potential package content conflict after content was moved from **openstack-tripleo-validations** to another package. (BZ#1877688)

Changes to the puppet-cinder component:

- This release adds support for the Dell EMC PowerStore Cinder Backend Driver. (BZ#1862545)

Changes to the puppet-tripleo component:

- This enhancement adds a new driver for the Dell EMC PowerStore to support Block Storage service back end servers. (BZ#1862546)

- This update fixes incorrect parameter names in Dell EMC Storage Templates. (BZ#1868620)

Changes to the python-networking-ovn component:

- Transmission of jumbo UDP frames on ML2/OVN routers depends on a kernel release that is not yet avaialbe.
  After receiving a jumbo UDP frame that exceeds the maximum transmission unit of the external network, ML2/OVN routers can return ICMP "fragmentation needed" packets back to the sending VM, where the sending application can break the payload into smaller packets. To determine the packet size, this feature depends on discovery of MTU limits along the south-to-north path.

  South-to-north path MTU discovery requires kernel-4.18.0-193.20.1.el8_2, which is scheduled for availability in a future release. To track availability of the kernel version, see https://bugzilla.redhat.com/show_bug.cgi?id=1860169. (BZ#1547074)

Changes to the python-os-brick component:

- This update modifies **get_device_info** to use lsscsi to get **[H:C:T:L]** values, making it possible to support more than 255 logical unit numbers (LUNs) and host logical unit (HLU) ID values. Previously, **get_device_info** used sg_scan to get these values, with a limit of 255.

  You can get two device types with **get_device_info**:

  - o /dev/disk/by-path/xxx, which is a symlink to /dev/sdX

  - o /dev/sdX
    sg_scan can process any device name, but lsscsi only shows /dev/sdx names.

    If the device is a symlink, **get_device_info** uses the device name that the device links to. Otherwise **get_device_info** uses the device name directly.

    Then **get_device_info** gets the device info '[H:C:T:L]' by comparing the device name with the last column of lsscsi output. (BZ#1872211)

- This update fixes an incompatibility that caused VxFlex volume detachment attempts to fail.
  A recent change in VxFlex cinder volume credentialing methods was not backward compatible with pre-existing volume attachments. If a VxFlex volume attachment was made before the credentialing method change, attempts to detach the volume failed.

  Now the detachments do not fail. (BZ#1869346)

Changes to the python-tripleoclient component:

- The entry in **/etc/hosts** for the undercloud duplicates anytime the Compute stack is updated on the undercloud and overcloud nodes. This occurs for split-stack deployments where the Controllers and Compute nodes are divided into multiple stacks.
  Other indications of this problem are the following:

  - mysql reporting errors about packets exceeding their maximum size.

  - The Orchestration service (heat) warning that templates are exceeding their maximum size.

  - The Workflow service (mistral) warning that fields are exceeding their maximum size. As a

workaround, in the file generated by running the **openstack overcloud export** command that is included in the Compute stack, under **ExtraHostFileEntries**, remove the erroneous entry for the undercloud. (BZ#1876153)

Changes to the tripleo-ansible component:

- This update increases the speed of stack updates in certain cases.
  Previously, stack update performance was degraded when the Ansible --limit option was not passed to ceph-ansible. During a stack update, ceph-ansible sometimes made idempotent updates on nodes even if the --limit argument was used.

  Now director intercepts the Ansible --limit option and passes it to the ceph-ansible excecution. The --limit option passed to commands starting with 'openstack overcloud' deploy is passed to the ceph-ansible execution to reduce the time required for stack updates.

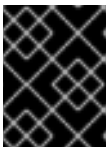

IMPORTANT

Always include the undercloud in the limit list when using this feature with ceph-ansible. (BZ#1855112)

4.5. RHBA-2021:0817 – RED HAT OPENSTACK PLATFORM 16.1.4 DIRECTOR BUG FIX ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2021:0817. Further information about this advisory is available at link: https://access.redhat.com/errata/RHBA-2021:0817.html.

Changes to the openstack-cinder component:

- Before this update, cloned encrypted volumes were inaccessible when using the Block Storage (cinder) service with the Key Manager (barbican) service. With this update, cloned encrypted volumes are now accessible when using the Block Storage service with the Key Manager service. (BZ#1889228)

- The 'all_tenants' key passed with a volume transfer request is removed because the database is unable to parse it. Removing this key allows the user to show the detail of a specific volume transfer by using the transfer name. Before this update, the 'all_tenants' key was removed only for admin users, which meant that non-admin users were unable to show volume transfers by using the transfer name. With this update, the 'all_tenants' key is now also removed for non-admins, allowing non-admins to show volume transfers by using the transfer name. (BZ#1847907)

- Before this update, the Block Storage (cinder) NEC back end driver occasionally returned invalid data when initializing a volume connection, which could cause live migration to fail. With this update, the NEC driver has been fixed to reliably return valid connection data. Live migration no longer fails due to invalid volume connection data. (BZ#1910854)

- Before this update, the Block Storage (cinder) service would always assign newly created volumes with the default volume type, even when the volume was created from another source, such as an image, snapshot or another volume. This resulted in volumes created from another source having a different volume type from the volume type of the source.
  With this update, the default volume type is assigned only after determining whether it should be assigned based on the volume type of the source. The volume type of volumes created from another source now match the volume type of the source. (BZ#1921735)

- Before this update, the **--server** option was being ignored when passed with the **cinder**

service-get-log command, which resulted in the logs for all hosts being returned instead of just the logs for a specific host. With this update, using the **--server** option correctly filters the logs for the specified host. (BZ#1728142)

Changes to the openstack-tripleo-common component:

- The **virt-admin** tool is now available for you to use to capture logs for reporting RHOSP bugs. This tool is useful for troubleshooting all libvirt and QEMU problems, as the logs provide the communications between libvirt and QEMU on the Compute nodes. You can use **virt-admin** to set the libvirt and QEMU debug log filters dynamically, without having to restart the **nova_libvirt** container.
  Perform the following steps to enable libvirt and QEMU log filters on a Compute node:

  1. Log in to the **nova_libvirt** container on the Compute node:

     ```
     $ sudo podman exec -it nova_libvirt /bin/bash
     ```

  2. Specify the name and location of the log file to send **virt-admin** output to:

     ```
     $ virt-admin daemon-log-outputs "1:file:/var/log/libvirt/libvirtd.log"
     ```

  3. Configure the filters you want to collect logs for:

     ```
     $ virt-admin daemon-log-filters \
      "1:libvirt 1:qemu 1:conf 1:security 3:event 3:json 3:file 3:object 1:util"
     ```

     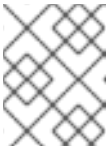

     NOTE

     When debugging issues with live migration, you must configure these filters on all source and destination Compute nodes.

  4. Repeat your test. After debugging is complete, upload the **libvirtd.log** to a bug.

  5. Disable the libvirt and QEMU log filters on the Compute nodes:

     ```
     $ virt-admin daemon-log-filters ""
     ```

  6. To confirm that the filters are removed, enter the following command:

     ```
     $ virt-admin daemon-log-filters
     ```

     This command returns an empty list when you have successfully removed the filters.

(BZ#1870199)

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

- Before this update, in-place upgrades from Red Hat OpenStack Platform 13 to 16.1 in a TLS everywhere environment used an incorrect rabbitmq password for the novajoin container. This caused the novajoin container on the undercloud to function incorrectly, which caused any overcloud node that ran an upgrade to fail with the following error:

  ```
  2020-11-24 20:01:31.569 7 ERROR join  File "/usr/lib/python3.6/site-
  packages/amqp/connection.py", line 639, in _on_close
  ```

```
2020-11-24 20:01:31.569 7 ERROR join     (class_id, method_id), ConnectionError)
2020-11-24 20:01:31.569 7 ERROR join amqp.exceptions.AccessRefused: (0, 0): (403)
ACCESS_REFUSED - Login was refused using authentication mechanism AMQPLAIN. For
detail see the broker logfile.
```

With this update, the upgrade from RHOSP 13 to 16.1 uses the correct rabbitmq password in a TLS everywhere environment so that the framework for upgrades can complete successfully. (BZ#1901157)

- With this enhancement, you can deploy the Red Hat Ceph Storage (RHCS) Dashboard on edge sites in a distributed compute node (DCN) architecture. (BZ#1793595)

- 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#1834185)

- There is currently a known issue with the mechanism that ensures the subscribed environments have the right DNF module stream set. The Advanced Virtualization repository is not always available in the subscription that the Ceph nodes use, which causes the upgrade or update of a Ceph node to fail when you try to enable virt:8.2.
  Workaround:

  Override the **DnfStreams** parameter in the upgrade or update environment file to prevent the Ceph upgrade from failing:

  ```
  parameter_defaults:
    ...
    DnfStreams: [{'module':'container-tools', 'stream':'2.0'}]
  ```

  

  NOTE

  The Advanced Virtualization DNF stream is not enforced when you use this workaround.

  For more information, see https://bugzilla.redhat.com/show_bug.cgi?id=1923887. (BZ#1866479)

- This enhancement adds support for heterogeneous storage configurations at the edge. Operators can now deploy edge sites with storage and sites without storage within the same DCN deployment. (BZ#1882058)

- The Block Storage backup service sometimes needs access to files on the host that would otherwise not be available in the container running the service. This enhancement adds the **CinderBackupOptVolumes** parameter, which you can use to specify additional container volume mounts for the Block Storage backup service. (BZ#1891828)

- Before this update, TLS-E on pre-provisioned nodes failed with the message: "--server cannot be used without providing --domain". With this update, the IDM domain name is detected by first resolving "ipa-ca" through DNS, then doing a reverse DNS lookup on the resulting IP

address. It might be necessary to add the PTR record, which is required for the reverse lookup, manually. (BZ#1874936)

- Before this update, you were required to use the **openstack overcloud external-upgrade run --tags online_upgrade** command to perform online database updates when upgrading from RHOSP 15 to RHOSP 16.1. With this update, you can now use the **openstack overcloud external-update run --tags online_upgrade** command. (BZ#1884556)

- Before this update, if you had **NovaComputeEnableKsm** enabled and you were using Red Hat Subscription Management to register the overcloud Compute nodes, the **qemu-kvm-common** package failed to install. This was because the configuration was sometimes applied before the Compute nodes were registered to the required repositories.
  With this update, **NovaComputeEnableKsm** is enabled only after the Compute nodes are registered to the required repositories by using Red Hat Subscription Management, which ensures that the **qemu-kvm-common** package is successfully installed. (BZ#1895894)

- Before this update, the connection data created by an iSCSI/LVM Block Storage back end was not stored persistently, which resulted in volumes not being accessible after a reboot. With this update, the connection data is stored persistently, and the volumes are accessible after a system reboot. (BZ#1898484)

- Before this update, when deployed at an edge site the Image (glance) service was not configured to access the Key Manager (barbican) service running on the central site's control plane. This resulted in the Image services running on edge sites being unable to access encryption keys stored in the Key Manager service.
  With this update, Image services running on edge sites are now configured to access the encryption keys stored in the Key Manager service. (BZ#1899761)

Changes to the puppet-collectd component:

- With this enhancement, you can configure the format of the plugin instance for the collectd virt plugin by using the **ExtraConfig** parameter **collectd::plugin::virt::plugin_instance_format**. This allows more granular metadata to be exposed in the metrics label for virtual machine instances, such as on which host the instance is running. (BZ#1878191)

- Before this update, when you configured the **collectd::plugin::virt::hostname_format** parameter with multiple values, director wrapped the values in double quotes. This caused the virt plugin to fail to load. With this update, when configuring **collectd::plugin::virt::hostname_format**, director no longer wraps multiple values in double quotes. (BZ#1902142)

Changes to the python-network-runner component:

- Before this update, a rebase in python-network-runner from 0.1.7 to 0.2.2 in OSP 16.1.3 caused ML2 Networking using Ansible to no longer function. With this update, python-networking-ansible is reverted to 0.1.7, and Ansible networking returns to a functioning state. For more information, see https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.1/html/bare_metal_provisioning/ml2-networking-ansible. (BZ#1909795)

Changes to the python-networking-ovn component:

- 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#1575512)

- Before this update, the OVN mechanism driver did not correctly merge its agent list with those stored in the Networking (neutron) service database. With this update, the results from the OVN and Networking service database are merged before the API returns the result. (BZ#1828889)

- This enhancement adds support for vlan transparency in the ML2/OVN mechanism driver with vlan and geneve network type drivers.
  With vlan transparency, you can manage vlan tags by using instances on Networking (neutron) service networks. You can create vlan interfaces on an instance and use any vlan tag without affecting other networks. The Networking service is not aware of these vlan tags.

  NOTE

- When using vlan transparency on a vlan type network, the inner and outer ethertype of the packets is 802.1Q (0x8100).

- The ML2/OVN mechanism driver does not support vlan transparency on flat provider networks.

(BZ#1846019)

Changes to the python-os-brick component:

- Before this update, instances that were created on a RHOSP 13 environment with PowerFlex, VxFlex and ScaleIO volume attachments failed restarting after an upgrade to RHOSP 16.x. This was because the RHOSP 16.x Compute service uses a new PowerFlex driver connection property to access volume attachments, which is not present in the connection properties of volumes attached to instances running on a RHOSP 13 environment. With this update, the error is no longer thrown if this connection property is missing, and instances with PowerFlex volume attachments created on a RHOSP 13 environment continue to function correctly after upgrading to RHOSP 16.x.

Changes to the python-paunch component:

- Before this update, if a user configured the **ContainerImagePrepare** parameter to use a custom tag, such as 'tag: "latest"' or 'tag: "16.1"', instead of the standard 'tag_from_label: "{version}-{release}"', the containers did not update to the latest container images.
  With this update, the container images are always fetched anytime a user runs a deployment action, including updates, and the image ID is checked against the running container to see if it needs to be rebuilt to consume the latest image. Containers are now always refreshed during deployment actions and restarted if they are updated.

  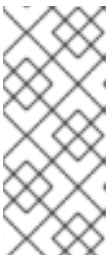

  NOTE

  This is a change from previous versions where the deployment checked only that the image existed rather than always fetching the image. If a user is reusing tags, for example, "latest", the containers might be updated on nodes if you perform actions such as scaling out. It is not recommended to use "latest" unless you are controlling container tags by using a Satellite server deployment.

  (BZ#1881476)

Changes to the python-tripleoclient component:

- Before this update, live migration failed when upgrading a TLS everywhere environment with local ephemeral storage and **UseTLSTransportForNbd** set to "False". This occurred because the default value of the **UseTLSTransportForNbd** configuration had changed from "False" in

RHOSP 13 to "True" in RHOSP 16.x, which resulted in the correct certifications not being included in the QEMU process containers.
With this update, director checks the configuration of the previously deployed environment for **global_config_settings** and uses it to ensure that the **UseTLSTransportForNbd** state stays the same in the upgrade as on previous deployment. If **global_config_settings** exists in the configuration file, then director checks the configuration of the **use_tls_for_nbd** key. If **global_config_settings** does not exist, the director evaluates the hieradata key **nova::compute::libvirt::qemu::nbd_tls**. Keeping the **UseTLSTransportForNbd** state the same in the upgraded deployment as on previous deployment ensures that live migration works. (BZ#1906698)

4.6. RHBA-2021:2097 – RED HAT OPENSTACK PLATFORM 16.1.6 DIRECTOR BUG FIX ADVISORY

Changes to the openstack-cinder component:

- In prior releases, the SolidFire driver created a duplicate volume whenever it retried an API request. This led to unexpected behavior due to the accumulation of unused volumes.
  With this update, the Block Storage service (cinder) checks for existing volume names before it creates a volume. When Block Storage service detects a read timeout, it immediately checks for volume creation to prevent invalid API calls. This update also adds the **sf_volume_create_timeout** option for the SolidFire driver so that you can set an appropriate timeout value for your environment. (BZ#1939398)

- This update fixes a bug that prevented **cinder list** from listing volumes when multiple filters were passed. (BZ#1843788)

- This update adds CHAP support to the Dell EMC PowerStore driver. (BZ#1905231)

- In prior releases, cinder NEC driver backups failed when the object was a snapshot. This occurred because the **snapshot** argument does not have the **volume_attachment** attribute. With this update, backups no longer refer to the **volume_attachment** attribute when the argument is **snapshot**. (BZ#1910855)

- This update fixes an issue that caused some API calls, such as create snapshot, to fail with an xNotPrimary error during workload re-balancing operations.
  When SolidFire is under heavy load or being upgraded, the SolidFire cluster might re-balance cluster workload by automatically moving connections from primary to secondary nodes. Previously, some API calls failed with an xNotPrimary error during these workload balance operations and were not retried.

  This update fixes the issue by adding the xNotPrimary exception to the SolidFire driver list of retryable exceptions. (BZ#1947474)

Changes to the openstack-heat component:

- This update makes it possible to use OS::Heat:Delay resources in heat templates. Previously, a variable naming conflict caused an assertion error during attempted completion of an OS::Heat::Delay resource. A variable was renamed to eliminate the conflict. (BZ#1868543)

Changes to the openstack-nova component:

- When an instance is created, the Compute (nova) service sanitizes the instance display name to generate a valid host name when DNS integration is enabled in the Networking (neutron) service.
  Before this update, the sanitization did not replace periods ('.') in instance names, for example,

'rhel-8.4'. This could result in display names being recognized as Fully Qualified Domain Names (FQDNs) which produced invalid host names. When instance names contained periods and DNS integration was enabled in the Networking service, the Networking service rejected the invalid host name, which resulted in a failure to create the instance and a HTTP 500 server error from the Compute service.

With this update, periods are now replaced by hyphens in instance names to prevent host names being parsed as FQDNs. You can continue to use free-form strings for instance display names. (BZ#1872314)

Changes to the openstack-tripleo-common component:

- This update modifies the registry metadata creator to handle containers with and without namespaces in their URI. On the undercloud you can now manage containers that comply with the following formats:
  undercloud_host:port/namespace/container:tag undercloud_host:port/container:tag

  Red Hat does not support more complex namespaces, such as undercloud_host:port/name/space/container:tag, when pushing to the undercloud. (BZ#1919445)

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

- After upgrading with the Leapp utility, Compute with OVS-DPDK workload does not function properly. Choose one of the following workaround options:
- remove /etc/modules-load.d/vfio-pci.conf, before compute upgrade
- restart compute ovs after compute upgrade. (BZ#1895887)
- This update fixes a configuration problem that caused Leapp upgrades to stop and fail while executing on a CephStorage node.
  Previously, CephStorage nodes were incorrectly configured to consume OpenStack highavailability, advanced-virt, and fast-datapath repos during Leapp upgrades.

  Now **UpgradeLeappCommand** options is configurable on a per-node basis, and uses the correct default for CephStorage nodes, and Leapp upgrades succeed for CephStorage nodes. (BZ#1936419)

Changes to the validations-common component:

- This update fixes a bug that caused failure of validations before **openstack undercloud upgrade** in some cases. Before this upgrade, a lack of permissions needed to access the requested logging directory sometimes resulted in the following failures:
  - Failure to log validation results
  - Failure of the validation run
  - Failure of artifacts collection from validation.
    This update adds a fallback logging directory. Validation results are logged and artifacts collected. (BZ#1895045)

4.7. RHBA-2021:3762 – RED HAT OPENSTACK PLATFORM 16.1.7 GENERAL AVAILABILITY ADVISORY

The bugs contained in this section are addressed by advisory RHSA-2021:3762. Further information about this advisory is available at link: https://access.redhat.com/errata/RHBA-2021:3762.html.

Changes to the diskimage-builder component:

- Before this update, the **appstream** and **baseos** repositories were always added to the repositories enabled by Red Hat Subscription Manager, with no way to override them. With this update, when you define the $REG_REPOS variable, no base repositories are added. You can control which repositories are added, but you must now include all repositories, including the equivalent repository for **baseos**, and **appstream** when required. (BZ#1906162)

Changes to the openstack-cinder component:

- Before this update, creating a volume from a snapshot of an encrypted volume could result in an unusable volume. When the destination volume is the same size as the source volume, creating an encrypted volume from a snapshot of an encrypted volume truncated the data in the new volume, which caused a size discrepancy.
  With this update, the RBD back end accounts for the encryption header and does not truncate the data so that creating a volume from a snapshot of an encrypted volume does not cause the error. (BZ#1987104)

- In previous releases, in Red Hat OpenStack Platform (RHOSP) deployments that use the Dell EMC XtremIO driver, attach volume operations waited for a timeout if iSCSI or FC targets were not connected to a RHOSP host. This caused attach volume operations to fail.
  This release adds port filtering support for the Dell EMC XtremIO driver to allow iSCSI or FC ports that are not in use to be ignored. (BZ#1930255)

- In previous releases, if Dell EMC PowerStore ports were configured for multiple purposes, such as iSCSI, replication, incorrect REST filtering caused the cinder driver to report that no accessible iSCSI targets were found.
  This release fixes the Dell EMC PowerStore REST filter functionality. (BZ#1945306)

- Before this update, a failure occurred when users wanted to delete the ***DEFAULT*** volume type. With this update, you can delete the ***DEFAULT*** volume type when it is not set as the value of the **default_volume_type** parameter in the **cinder.conf** file. The default value of the **default_volume_type** parameter is ***DEFAULT*** so you must set it to an appropriate volume type, for example 'tripleo', so that you can delete the ***DEFAULT*** volume type. (BZ#1947415)

Changes to the openstack-manila-ui component:

- Before this update, the Shared File Systems service (manila) dashboard had dynamic form elements whose names could potentially cause the forms to become unresponsive. This meant that the creation of share groups, share networks, and shares within share networks did not function.
  With this update, dynamic elements whose names might be problematic are encoded. The creation of share groups, share networks, and shares within share networks functions normally. (BZ#1938212)

Changes to the openstack-neutron component:

- The logic to detect the hypervisor hostname has been fixed and now returns the result consistent with **libvirt** driver in the Compute service (nova). With this fix, you no longer need to specify the **resource_provider_hypervisors** option when you use the guaranteed minimum bandwidth QoS feature.
  With this update, a new option, **resource_provider_default_hypervisor**, has been added to the Modular Layer 2 with the Open Virtual Network mechanism driver (ML2/OVN) to replace the default hypervisor name. The option locates the root resource provider without giving a

complete list of interfaces or bridges in the **resource_provider_hypervisors** option in case it has to be customized by the user. This new option is located in the **[ovs]** ini-section for the **ovs-agent**, and in the **[sriov_nic]** ini-section for the **sriov-agent**. (BZ#1900500)

Changes to the openstack-octavia component:

- With this update, there is resolution to a problem that prevented the RHOSP Load-balancing service (octavia) to fail over load balancers with multiple failed amphorae. (BZ#1974831)

- Before this update, when a configuration change to a Load-balancing service amphora caused an haproxy reload, the process consumed a lot of memory that could lead to memory allocation errors. The issue was caused by the **lo** interface not being configured in the amphora-haproxy namespace in the amphora. With this update, the namespace issue has been corrected and the problem is resolved. (BZ#1975790)

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

- Before this update, upgrading a Red Hat OpenStack Platform (RHOSP) 13 environment that has been deployed with ML2-OVN, to RHOSP 16.1 caused the upgrade process to fail on the Controller nodes due to an SELinux denial issue. With this update, the correct SELinux label is applied to OVN and resolves the issue. For more information, see the Red Hat Knowledgebase solution OVN fails to configure after reboot during OSP-13 → OSP-16.1 FFU . (BZ#1997351)

- Before this update, if your environment was deployed with a TLS-Everywhere architecture and it used the deprecated **authconfig** utility to configure authentication on your system, you had to configure your RHEL 8 system with the **authselect** utility. Without performing this action, the **leapp** process failed with the inhibitor named **Missing required answers in the answer file**. The workaround was to add **sudo leapp answer --section authselect_check.confirm=True --add** in the **LeappInitCommand** in the upgrades environment file. With this update, the configuration entry is no longer needed, and the upgrade now completes without intervention. (BZ#1952574)

- Before this update, the Red Hat Enterprise Linux (RHEL) in-place upgrade tool, LEAPP, stalled because it encountered loaded kernel modules that are no longer provided in RHEL 8. Also, LEAPP upgraded RHEL to a version that is not supported by Red Hat OpenStack Platform (RHOSP). With this update, the manual configurations that you had to perform to workaround these two issues are no longer required. (For more information, see BZ1962365. (BZ#1962365)

- With this update, the memory limit for the **collectd** container has been increased to 512 MB. When this limit is exceeded, the container restarts. (BZ#1969895)

- Before this update, removal of the **python2** packages for the Red Hat Enterprise Linux (RHEL) in-place upgrade tool, LEAPP, was unsuccessful. This failure was caused by a DNF **exclude** option that retained the LEAPP packages. With this update, automation has now been included to ensure that the necessary LEAPP packages are successfully removed. (BZ#2008976)

- Before this update, an upgradable **mariadb-server** package in the RHEL repository caused the package manager to upgrade the **mariadb-server** package on the host, which interfered with the containerized **mariadb-server** that pre-exists on the same host. With this update, the Red Hat OpenStack Platform (RHOSP) director removes the **mariadb-server** package from any hosts that also have the containerized MariaDB, and the RHOSP FFU process continues. (BZ#2015325)

- This enhancement adds the new **CinderRpcResponseTimeout** and **CinderApiWsgiTimeout** parameters to support tuning RPC and API WSGI timeouts in the Block Storage service (cinder). Default timeout values might not be adequate for large deployments and in situations where transactions might be delayed due to system load.

It is now possible to tune the RPC and API WSGI timeouts to prevent transactions prematurely timing out. (BZ#1930806)

Changes to the puppet-collectd component:

- Previously, the **PluginInstanceFormat** parameter for collectd accepted only one of the following values: 'none', 'name', 'uuid', or 'metadata'. With this update, you can now specify more than one value for the **PluginInstanceFormat** parameter, resulting in more information being sent in the **plugin_instance** label of collectd metrics. (BZ#1956887)

Changes to the python-networking-ovn component:

- Currently, there is a known issue where it is not possible to simulate certain real-life scenarios when the MAC-IP addresses of a port are unknown. The RHOSP Networking service (neutron) directly specifies the MAC-IP of a port even if DHCP or security groups are not configured. Workaround: upgrade to RHOSP 16.1.7 and install ML2/OVN v21.03. If DHCP and port security are disabled, then the addresses field of a port does not include its MAC-IP address pairs, and ML2/OVN can use the MAC learning capabilities to send traffic only to the desired port. (BZ#1898198)

Changes to the python-os-brick component:

- Before this update, there were unhandled exceptions during connection to iSCSI portals. For example, failures in **iscsiadm -m session**. This occurred because the **_connect_vol** threads can abort unexpectedly in some failure patterns, and this abort causes a hang in subsequent steps while waiting for results from **_connct_vol** threads.
  With this update, any exceptions during connection to iSCSI portals are handled in the **_connect_vol** method correctly and avoids any unexpected abort without updating thread results. (BZ#1977792)

Changes to the python-tripleoclient component:

- With this update, the **tripleo validator** command now accepts variables and environment variables in a key-value pair format. In past releases, only JSON dictionaries allowed environment variables.

  ```
  openstack tripleo validator run \
  [--extra-vars key1=<val1>[,key2=val2 --extra-vars key3=<val3>] \
  | --extra-vars-file EXTRA_VARS_FILE] \
  [--extra-env-vars key1=<val1>[,key2=val2 --extra-env-vars key3=<val3>]]
  (--validation <validation_id>[,<validation_id>,...] | --group <group>[,<group>,...])
  ```

  Example

  ```
  $ openstack tripleo validator run --validation check-cpu,check-ram --extra-vars
  minimal_ram_gb=8 --extra-vars minimal_cpu_count=2
  ```

  For the complete list of supported options, run:

  ```
  $ openstack tripleo validator run --help
  ```

  (BZ#1959492)

- Before this update, during a tripleo validation on an OpenStack component, the following exception error occurred:

```
Unhandled exception during validation run.
```

This error occurred because a variable in the code was referenced, but never assigned.

With this update, this problem has been fixed and validations run without this error. (BZ#1959866)

Changes to the tripleo-ansible component:

- Before this update, an optional feature of the RHOSP Load-balancing service (octavia), log offloading, was not correctly configured during deployment. As a result of this problem, the Load-balancing service was not receiving logs from the amphorae. This update resolves the issue. (BZ#1981652)

- Before this update, changes to **KernelArgs** parameters caused errors in the Red Hat OpenStack Platform (RHOSP) fast forward upgrade (FFU) process for version 13 to version 16:

- Duplicate entries appeared in **/etc/default/grub**.

- Duplicate entries appeared in the kernel command line.

- Nodes rebooted during the RHOSP upgrade.
  These errors were caused when the **KernelArgs** parameter, or the order of values in the string, changed or when a **KernelArgs** parameter was added.

  With this update, TripleO has added upgrade tasks in **kernel-boot-params-baremetal-ansible.yaml** to migrate from **TRIPLEO_HEAT_TEMPLATE_KERNEL_ARGS** to **GRUB_TRIPLEO_HEAT_TEMPLATE_KERNEL_ARGS**.

  This change was made to accommodate the Red Hat Enterprise Linux (RHEL) in-place upgrade tool, LEAPP, which is used to upgrade RHEL from version 7 to version 8, during the RHOSP version 13 to version 16 FFU process. LEAPP understands GRUB parameters only when the parameters start with **GRUB_** in **/etc/default/grub**.

  Despite this update, you must manually inspect each **KernelArgs** value to ensure that it matches the value for all hosts in the corresponding role.

  The **KernelArgs** value may come from the **PreNetworkConfig** implementation from either the default tripleo-heat-templates or third-party heat templates.

  If you find any mismatches, change the value of the **KernelArgs** parameter in the corresponding role to match the value of **KernelArgs** on the hosts. Perform these checks before running the **openstack overcloud upgrade prepare** command.

  You can use the following script to check **KernelArgs** values:

```
tripleo-ansible-inventory --static-yaml-inventory inventory.yaml
KernelArgs='< KernelArgs_ FROM_THT >'
ansible -i inventory.yaml ComputeSriov -m shell -b -a "cat /proc/cmdline | grep
'${KernelArgs}'"
```

  (BZ#1980829)

4.8. RHBA-2022:0986 – RED HAT OPENSTACK PLATFORM 16.1.8 BUG FIX AND ENHANCEMENT ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2022:0986. For more information, see link:https://access.redhat.com/errata/RHBA-2022:0986.html.

+ Changes to the openstack-cinder component:

- Before this update, the GPFS SpectrumScale driver in the Block Storage service (cinder) did not correctly detect whether the storage back end supported copy-on-write (COW) mode. As a result, the driver disabled COW features, such as the ability to rapidly create volumes from an image. Sometimes, this caused some instances to time out when booting multiple instances simultaneously from an image.
  With this update, the GPFS SpectrumScale driver properly detects COW support for storage back ends. (BZ#1960639)

- Before this update, when creating a snapshot with PowerMaxOS 5978.711, REST experienced a payload response change and caused the device label to modify its format. The underlying data from the solutions enabler changed, and no longer contained a colon character (:). This resulted in an IndexError exception in the PowerMax Driver:

  ```
  IndexError: list index out of range
  ```

  With this update, the problem is resolved in PowerMaxOS 5978.711 and later. (BZ#1992159)

- Before this update, the OpenStack NFS driver blocked attempts to delete snapshots in an error state when snapshot support is disabled. New or existing snapshots are placed in an error state when snapshot support is disabled, but users could not remove these failed snapshots. With this update, users can now remove NFS snapshots in error status. (BZ#1741453)

- Before this update, the PowerMax driver used a mechanism for storing and maintaining information on shared volume connections that did not work with previously created legacy volumes. This caused live migration to fail for volumes that were created before the PowerMax migration code was introduced. Now, the PowerMax live migration code is updated to work with legacy volumes so that live migrations do not fail. (BZ#1987957)

- This update fixes a bug that omitted details from the output of the **openstack volume backup list** command when the output exceeded 1000 lines. (BZ#1999634)

Changes to the openstack-tripleo-common component:

- With this update, the telemetry healthchecks have been made more robust and the way the healthchecks are parsed has been simplified.
  To get verbose mode when you run the healthcheck directly, run the command **sudo podman -u root -e "HEALTHCHECK_DEBUG=1" <container> /openstack/healthcheck** (BZ#1910939)

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

- As of this release, the Red Hat supported method of updating OVN is aligned to the upstream OVN updgrade steps. (BZ#2052411)

- Before this release, the collectd container failed to start on Compute nodes because a dpdk-telemetry collectd configuration file was being automatically created despite there being no dpdk-telemetry plugin installed.
  As of this release, dpdk_telemetry configuration files have been removed from the the collectd container. (BZ#1996865)

- Enable the experimental **rsyslog reopenOnTruncate** to ensure that rsyslog immediately recognizes when a logrotation happens on a file. The setting affects every service configured to work with rsyslog.
  With **rsyslog reopenOnTruncate** disabled, rsyslog waits for a log file to fill to its original capacity before consuming any additional logs. (BZ#1939964)

- With this update the **CollectdContainerAdditionalCapAdd** variable is added to the deployment tool. This variable is a comma separated list of additional collectd container capabilities. (BZ#1984095)

- With this update, the **LeapActorsToRemove** heat parameter is introduced so that you can remove specific actors from the leapp process if those actors inhibit the upgrade. The **LeapActorsToRemove** heat parameter is role-specific for flexibility. (BZ#1984873)

Changes to the puppet-tripleo component:

- This enhancement prepares your environment for update of the metrics_qdr service to a newer AMQ Interconnect release, which requires import of the CA certificate contents from the Service Telemetry Framework (STF) deployment. Changes are not yet required by administrators when deploying or updating Red Hat OpenStack Service Platform (RHOSP) as the metrics_qdr service has not yet been updated. This functionality is in preparation of the metrics_qdr service update in a future release.
  The following procedure will be required once https://bugzilla.redhat.com/show_bug.cgi?id=1949169 has shipped.

  This update corrects this problem by providing a new Orchestration service (heat) parameter, **MetricsQdrSSLProfiles**.

  To obtain a Red Hat OpenShift TLS certificate, run the following commands:

  ```
  $ oc get secrets
  $ oc get secret/default-interconnect-selfsigned -o jsonpath='{.data.ca\.crt}' | base64 -d
  ```

  Add the **MetricsQdrSSLProfiles** parameter with the contents of your Red Hat OpenShift TLS certificate to a custom environment file:

  ```
  MetricsQdrSSLProfiles:
    - name: sslProfile
      caCertFileContent: |
        -----BEGIN CERTIFICATE-----
        ...
        TOpbgNIPcz0sIoNK3Be0jUcYHVMPKGMR2kk=
        -----END CERTIFICATE-----
  ```

  Then, redeploy your overcloud with the **openstack overcloud deploy** command. (BZ#1949168)

- This update corrects an error that prevented the proper use of the Cinder **powermax_port_groups** parameter. (BZ#2029608)

Changes to the python-os-brick component:

- Before this update, os-brick did not include a **[global]** section to contain the options it sets in a temporary configuration file, which is a requirement with Octopus (release 15.2.0+). As a result, connection information could not be found when using os-brick and a Ceph Octopus or later

client, and a connection to the Ceph storage backend could not be established. Now, the connection options are included under a '[global]' section in the temporary configuration file. This fix is backward compatible to the Hammer release (0.94.0+) of Ceph. (BZ#2023413)

4.9. RHBA-2022:8795 – RED HAT OPENSTACK PLATFORM 16.1.9 BUG FIX AND ENHANCEMENT ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2022:8795. For more information, see link: https://access.redhat.com/errata/RHBA-2022:8795.html.

Changes to the openstack-cinder component:

- Before this update, a race condition occurred when the Compute service (nova) requested the Block Storage service (cinder) to detach a volume and there was an external request to delete the volume. The race condition resulted in the volume failing to detach, the volume being deleted, and the Compute service being unable to remove the non-existent volume. With this update, the race condition is resolved. (BZ#1977322)

- Before this update, if you imported a backup record for a backup ID that currently existed, the import operation would correctly fail, but the existing backup record would incorrectly be deleted. With this update, the existing backup record is not deleted under this scenario. (BZ#1802263)

- Before this update, NetApp ONTAP Block Storage (cinder) driver QoS policy groups were deleted when the associated volume was moved. With this update, QoS policy groups are associated permanently to the LUN or file that represents the volume. (BZ#1951485)

- Before this update, a do_sync_check operation could result in the incorrect deletion of non-temporary snapshots from a volume because there was no check for non-temporary snapshots deletion during the do_sync_check operation. With this update, there is a check to determine if a snapshot must be deleted. The do_sync_check operation does not perform unnecessary non-temporary snapshot deletions.
  Before this update, there was a case mismatch in the conditional while checking if a storage group was a child of a parent storage group. While modifying the storage group, errors indicated that the parent storage group already contained the child storage group. With this update, the patterns used in the conditional are not case-sensitive and you can modify the storage group successfully. (BZ#2129310)

Changes to the openstack-ironic component:

- Before this update, if there were repeated transient connectivity issues between the ironic-conductor service and a remote Baseboard Management Controller (BMC) using the Redfish hardware type when session authentication was used, the intermittent loss of connectivity could collide with a point where authentication was retried due to the in-memory credentials expiring. If this collision occurred, there was a loss of overall connectivity, which persisted due to the internal session cache built into the openstack-ironic-conductor service. With this update, support to detect and renegotiate in cases of this error were added to the Python DMTF Redfish library, sushy, and the openstack-ironic service. Intermittent connectivity failures colliding with session credential re-authentication no longer results in a complete loss of ability to communicate with the BMC until the openstack-ironic-conductor service is restarted. (BZ#2027544)

Changes to the openstack-manila component:

- Before this update, the API that the Shared File Systems service (manila) uses to provision storage on NetApp ONTAP All Flash Fabric-Attached (AFF) storage systems caused Shared

File Systems service shares to be thinly provisioned. The API did not enforce space guarantees, even when requested through the Shared File Systems service share type. With this update, the driver sets appropriate parameters for the NetApp ONTAP 9 API to work with AFF storage as well as traditional FAS storage systems. The API enforces space guarantees on NetApp ONTAP storage through the Shared File Systems service share types. (BZ#1968228)

Changes to the openstack-nova component:

- There is currently a known issue when live migrating instances that have CPUs that are incompatible with the destination host CPUs.
  Workaround: Add the following configuration in the **nova.conf** file of each affected Compute node to skip CPU comparison on the destination host:

  ```
  [workarounds]
  skip_cpu_compare_on_dest = True
  ```

  (BZ#2076884)

- Before this update, block device mapping updates by the libvirt driver on the destination host were not persisted during live migration. With specific storage back ends or configurations, for example, when using the **n[workarounds]/rbd_volume_local_attach=True** config option, certain operations on volume attachments, for example detaching, after a live migration did not work. With this update, you can correctly persist any block device mapping updates done by the libvirt driver on the destination host. Operations on affected volumes, such as detaching, succeed after live migration. (BZ#2089382)

Changes to the openstack-octavia component:

- Before this update, the Virtual IP (VIP) address of UDP-only load balancers in active-standby mode was not reachable. With this update, the issue is fixed. (BZ#2078377)

- Before this update, Conntrack was enabled in the Amphora VM for any type of packet, but it is only required for the User Datagram Protocol (UDP) and Stream Control Transmission Protocol (SCTP). With this update, Conntrack is now disabled for Transmission Control Protocol (TCP) flows, preventing some performance issues when a user generates a lot of connections that fill the Conntrack table. (BZ#2123225)

- Before this update, members in the ERROR operating status might have been updated briefly to ONLINE during a Load Balancer configuration change. With this update, the issue is fixed. (BZ#1996756)

- Before this update, the provisioning status of a load balancer was set to ERROR too early when an error occurred, making the load balancer mutable before the execution of the tasks for these resources was finished. With this update, the issue is fixed. (BZ#2040697)

- Before this update, a SELinux issue triggered errors when using the ICMP monitor in the Load-balancing service (octavia) amphora driver. With this update, the SELinux issue is fixed. (BZ#2096387)

Changes to the openstack-tripleo-common component:

- RHSA-2022:6969 introduced the process to clean up files in the /var/lib/mistral directory in the undercloud but the process consistently failed when the Load-balancing service (octavia) or Red Hat Ceph Storage was enabled because these services created additional directories, which the cleanup process could not properly remove. Some deployment actions, such as scale out, consistently failed if the Load-balancing service or Ceph Storage was enabled. With this

update, Mistral no longer executes the cleanup. Users must manually delete files if they want to enforce the reduced permission of the files in the /var/lib/mistral directory. Deployment actions no longer fail because of a permission error. (BZ#2138184)

Changes to the puppet-rsyslog component:

- With this update, the Rsyslog environment configuration supports an array of Elasticsearch targets. In previous releases, you could only specify a single target. You can now specify multiple Elasticsearch targets as a list of endpoints to send logs. (BZ#1945334)

Changes to the python-dogpile-cache component:

- Before this update, dogpile.cache support for dead_retry and socket_timeout was not implemented for the memcached back end. The oslo.cache mechanism filled the arguments dictionary with values for dead_retry and socket_timeout, but dogpile.cache ignored the values so the defaults of 30s for dead_retry and 3s for socket_timeout were used. When using dogpile.cache.memcached as the cache back end on the Identity service (keystone), and then taking down one of the memcached instances, the memcache server objects set their deaduntil value to 30 seconds in the future. When a request came in to an API server with two memcached servers configured, one of which was unroutable, it took approximately 15 seconds for it to try each of those servers in each thread it created and reach the three-second socket timeout limit every time it encountered the one that was down. By the time the user issued another request, the deaduntil value was reached and the whole cycle was repeated. With this update, dogpile.cache consumes dead_retry and socket_timeout arguments passed by oslo.cache. (BZ#2100879)

Changes to the python-networking-ovn component:

- This update in RHOSP 16.1.9 fixes a bug that causes the Networking service (neutron) to fail to start after an update to RHOSP 16.1.8 and also causes OVN database instability after updates to RHOSP 16.1.8.
  Instead of updating to RHOSP 16.1.8, update directly to RHOSP 16.1.9. (BZ#2125824)

- With this update, you can now migrate an ML2/OVS deployment with the iptables_hybrid firewall driver to ML2/OVN. (BZ#2022040)

- When a load balancer is created in a tenant network with a Virtual IP (VIP) and members, and the tenant network is connected to a router that is connected to the provider network, the Open Virtual Network (OVN) load balancer is associated with the OVN logical router. If the 'router' option was used for nat-addresses, ovn-controller sent GARP packets for that VIP on the provider network. As there was nothing to prevent different tenants in OpenStack from creating a subnet with the same Classless Inter-Domain Routing (CIDR) number and a load balancer with the same VIP, there could be several ovn-controllers generating GARP packets on the provider network for the same IP, each one with the MAC of the logical router port belonging to each tenant. This setup could be an issue for the physical network infrastructure. With this update, a new option (exclude-lb-vips-from-garp) is added in OVN[1] on the router gateway port. This flag ensures that no GARP packets are sent for the load balancer VIPs. (BZ#2064709)

- Before this update, it was possible to add members without stating which subnet they belonged to, but they should be in the same subnet as the Virtual IP (VIP) port. If the subnet of the members is different to the VIP subnet, the members are created but incorrectly configured because there is no connectivity to them. With this update, members without a subnet are only accepted if the IP of the member belongs to the Classless Inter-Domain Routing (CIDR) number of the VIP subnet, as that is the subnet associated to the load balancer used to obtain the subnet for the members that do not have it. Member creation without a subnet is rejected if its IP does not belong to the VIP subnet CIDR. (BZ#2122925)

Changes to the python-octaviaclient component:

- Before this update, python-octaviaclient did not display the full list of load balancers when the user had more than 1,000 load balancers. With this update, the OpenStack Load-balancing service (Octavia) displays all load balancers. (BZ#1996088) Changes to the python-openstackclient component of Bugzilla:

Changes to the tripleo-ansible component:

- Before this update, the Load-balancing services (octavia) were restarted many times during deployments or updates. With this update, the services are restarted only when required, preventing potential interruptions of the control plane. (BZ#2057604)

- Before this update, a nonexistent gateway address was configured on the load-balancing management network. This caused excessive Address Resolution Protocol (ARP) requests on the load-balancing management network. (BZ#1961162)

- With this update, the **port_security** parameter of the Load-balancing service (octavia) management network is now enabled. (BZ#1982268)