Red Hat OpenStack Platform 16.0

Release Notes

Release details for Red Hat OpenStack Platform 16.0

Last Updated: 2020-05-29

Red Hat OpenStack Platform 16.0 Release Notes

Release details for Red Hat OpenStack Platform 16.0

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

CHAPTER 1. INTRODUCTION 3
1.1. ABOUT THIS RELEASE 3
1.2. REQUIREMENTS 3
1.3. DEPLOYMENT LIMITS 3
1.4. DATABASE SIZE MANAGEMENT 4
1.5. CERTIFIED DRIVERS AND PLUG-INS 4
1.6. CERTIFIED GUEST OPERATING SYSTEMS 4
1.7. PRODUCT CERTIFICATION CATALOG 4
1.8. BARE METAL PROVISIONING OPERATING SYSTEMS 4
1.9. HYPERVISOR SUPPORT 4
1.10. CONTENT DELIVERY NETWORK (CDN) REPOSITORIES 4
1.10.1. Undercloud repositories 5
1.10.2. Overcloud repositories 6
1.11. PRODUCT SUPPORT 8

CHAPTER 2. TOP NEW FEATURES 9
2.1. COMPUTE 9
2.2. NETWORKING 9
2.3. STORAGE 9
2.4. CEPH STORAGE 10
2.5. CLOUD OPS 10
2.6. TECHNOLOGY PREVIEWS 11
2.6.1. New Technology Previews 11

CHAPTER 3. RELEASE INFORMATION 14
3.1. RED HAT OPENSTACK PLATFORM 16.0 GA 14
3.1.1. Bug Fix 14
3.1.2. Enhancements 14
3.1.3. Technology Preview 18
3.1.4. Release Notes 24
3.1.5. Known Issues 24
3.1.6. Removed Functionality 28
3.2. RED HAT OPENSTACK PLATFORM 16.0.1 MAINTENANCE RELEASE 4 MARCH 2020 29
3.2.1. Enhancements 29
3.2.2. Known Issues 29
3.3. RED HAT OPENSTACK PLATFORM 16.0.2 MAINTENANCE RELEASE 14 MAY 2020 32
3.3.1. Enhancements 32
3.3.2. Technology Preview 33
3.3.3. Release Notes 33
3.3.4. Known Issues 33

CHAPTER 4. TECHNICAL NOTES 35
4.1. RHEA-2020:0283 – RED HAT OPENSTACK PLATFORM 16.0 GENERAL AVAILABILITY ADVISORY 35
4.2. RHBA-2020:2114 – RED HAT OPENSTACK PLATFORM 16.0.2 ADVISORY 44

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 Red Hat Enterprise Linux version 8.1.

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

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

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

NOTE

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

1.3. DEPLOYMENT LIMITS

For a list of deployment limits for Red Hat OpenStack Platform, see Deployment Limits for Red Hat OpenStack Platform.

1.4. DATABASE SIZE MANAGEMENT

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

1.5. CERTIFIED DRIVERS AND PLUG-INS

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

1.6. CERTIFIED GUEST OPERATING SYSTEMS

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

1.7. PRODUCT CERTIFICATION CATALOG

For a list of the Red Hat Official Product Certification Catalog, see Product Certification Catalog.

1.8. BARE METAL PROVISIONING OPERATING SYSTEMS

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

1.9. HYPERVISOR SUPPORT

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

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

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

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

1.10. CONTENT DELIVERY NETWORK (CDN) REPOSITORIES

This section describes the repositories required to deploy Red Hat OpenStack Platform 16.0.

You can install Red Hat OpenStack Platform 16.0 through the Content Delivery Network (CDN) using **subscription-manager**. For more information, see Preparing the undercloud.

WARNING

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

1.10.1. Undercloud repositories

Red Hat OpenStack Platform 16.0 runs on Red Hat Enterprise Linux {rhelvernum}. Before enabling repositories, lock the director to a version with the **subscription-manager release** command:

```
$ sudo subscription-manager release --set={rhelvernum}
```

Enable the following repositories for the installation and configuration of the undercloud.

Core repositories

The following table lists core repositories for installing the undercloud.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for x86_64 - BaseOS (RPMs) Extended Update Support (EUS)	**rhel-8-for-x86_64-baseos-eus-rpms**	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 8 for x86_64 - AppStream (RPMs)	**rhel-8-for-x86_64-appstream-rpms**	Contains Red Hat OpenStack Platform dependencies.
Red Hat Enterprise Linux 8 for x86_64 - High Availability (RPMs) Extended Update Support (EUS)	**rhel-8-for-x86_64-highavailability-eus-rpms**	High availability tools for Red Hat Enterprise Linux. Used for Controller node high availability.
Red Hat Ansible Engine 2.8 for RHEL 8 x86_64 (RPMs)	**ansible-2.8-for-rhel-8-x86_64-rpms**	Ansible Engine for Red Hat Enterprise Linux. Used to provide the latest version of Ansible.
Red Hat Satellite Tools for RHEL 8 Server RPMs x86_64	**satellite-tools-6.5-for-rhel-8-x86_64-rpms**	Tools for managing hosts with Red Hat Satellite 6.
Red Hat OpenStack Platform 16.0 for RHEL 8 (RPMs)	**openstack-16-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.

IBM POWER repositories

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

Name	Repository	Description of requirement
Red Hat Enterprise Linux for IBM Power, little endian - BaseOS (RPMs)	**rhel-8-for-ppc64le-baseos-rpms**	Base operating system repository for ppc64le systems.
Red Hat Enterprise Linux 8 for IBM Power, little endian - AppStream (RPMs)	**rhel-8-for-ppc64le-appstream-rpms**	Contains Red Hat OpenStack Platform dependencies.
Red Hat Enterprise Linux 8 for IBM Power, little endian - High Availability (RPMs)	**rhel-8-for-ppc64le-highavailability-rpms**	High availability tools for Red Hat Enterprise Linux. Used for Controller node high availability.
Red Hat Ansible Engine 2.8 for RHEL 8 IBM Power, little endian (RPMs)	**ansible-2.8-for-rhel-8-ppc64le-rpms**	Ansible Engine for Red Hat Enterprise Linux. Provides the latest version of Ansible.
Red Hat OpenStack Platform 16.0 for RHEL 8 (RPMs)	**openstack-16-for-rhel-8-ppc64le-rpms**	Core Red Hat OpenStack Platform repository for ppc64le systems.

1.10.2. Overcloud repositories

Red Hat OpenStack Platform 16.0 runs on Red Hat Enterprise Linux {rhelvernum}. After overcloud deployment, lock the to a version with the **subscription-manager release** command on each node:

```
$ sudo subscription-manager release --set={rhelvernum}
```

You must enable the following repositories to install and configure the overcloud.

Core repositories

The following table lists core repositories for installing the overcloud.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for x86_64 - BaseOS (RPMs) Extended Update Support (EUS)	**rhel-8-for-x86_64-baseos-eus-rpms**	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 8 for x86_64 - AppStream (RPMs)	**rhel-8-for-x86_64-appstream-rpms**	Contains Red Hat OpenStack Platform dependencies.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for x86_64 - High Availability (RPMs) Extended Update Support (EUS)	**rhel-8-for-x86_64-highavailability-eus-rpms**	High availability tools for Red Hat Enterprise Linux. Used for Controller node high availability.
Red Hat Ansible Engine 2.8 for RHEL 8 x86_64 (RPMs)	**ansible-2.8-for-rhel-8-x86_64-rpms**	Ansible Engine for Red Hat Enterprise Linux. Used to provide the latest version of Ansible.
Advanced Virtualization for RHEL 8 x86_64 (RPMs)	**advanced-virt-for-rhel-8-x86_64-rpms**	Provides virtualization packages for OpenStack Platform.
Red Hat Satellite Tools for RHEL 8 Server RPMs x86_64	**satellite-tools-6.5-for-rhel-8-x86_64-rpms**	Tools for managing hosts with Red Hat Satellite 6.
Red Hat OpenStack Platform 16.0 for RHEL 8 (RPMs)	**openstack-16-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.

Real Time repositories

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

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

IBM POWER repositories

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

Name	Repository	Description of requirement
Red Hat Enterprise Linux for IBM Power, little endian - BaseOS (RPMs)	**rhel-8-for-ppc64le-baseos-rpms**	Base operating system repository for ppc64le systems.
Red Hat Enterprise Linux 8 for IBM Power, little endian - AppStream (RPMs)	**rhel-8-for-ppc64le-appstream-rpms**	Contains Red Hat OpenStack Platform dependencies.
Red Hat Enterprise Linux 8 for IBM Power, little endian - High Availability (RPMs)	**rhel-8-for-ppc64le-highavailability-rpms**	High availability tools for Red Hat Enterprise Linux. Used for Controller node high availability.
Red Hat Ansible Engine 2.8 for RHEL 8 IBM Power, little endian (RPMs)	**ansible-2.8-for-rhel-8-ppc64le-rpms**	Ansible Engine for Red Hat Enterprise Linux. Used to provide the latest version of Ansible.
Red Hat OpenStack Platform 16.0 for RHEL 8 (RPMs)	**openstack-16-for-rhel-8-ppc64le-rpms**	Core Red Hat OpenStack Platform repository for ppc64le systems.

1.11. PRODUCT SUPPORT

Available resources include:

Customer Portal

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

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

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

Mailing Lists

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

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

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

CHAPTER 2. TOP NEW FEATURES

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

2.1. COMPUTE

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

Multi-cell deployments using Cells v2

OpenStack Compute is now powered by cells by default. You can configure larger deployments to use multiple cells, each cell with specific Compute nodes and database. Global services control placement and fail-safe operations, and the separation into cells help improve security and process isolation.

Colocation of pinned and floating instances on a single host

You can now schedule instances with pinned CPUs (**hw:cpu_policy=dedicated**) on the same host as instances that use floating CPUs (**hw:cpu_policy=shared**). It is no longer necessary to use host aggregates to ensure these instance types run on separate hosts.

Live migration for instances with SR-IOV (Single root I/O virtualization)

Instances configured with SR-IOV interfaces can now be live migrated. For direct mode SR-IOV interfaces, this operation will incur some network downtime as the interface is detached and reattached as part of the migration. This is not an issue for indirect mode SR-IOV interfaces.

Live migration of pinned instances

This NUMA-aware live migration feature allows you to live migrate instances with a NUMA topology.

Bandwidth-aware scheduling

You can now create instances that request a guaranteed minimum bandwidth using a Quality of Service (QoS) policy. The Compute scheduling service selects a host for the instance that satisfies the guaranteed minimum bandwidth request.

2.2. NETWORKING

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

ACL support for Load-balancing service (Octavia)

The Red Hat OpenStack Platform Load-balancing service (Octavia) now supports VIP access control lists (ACL) to limit incoming traffic to a listener to a set of allowed source IP addresses (CIDRs).

OVN internal API TLS/SSL support

Red Hat OpenStack Platform now supports the encryption of internal API traffic for OVN using Transport Layer Security (TLS).

OVN deployment over IPv6

Red Hat OpenStack Platform now supports deploying OVN on an IPv6 network.

2.3. STORAGE

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

Block Storage service changes keys when cloning volumes

With this release, when the Block Storage service (cinder) clones encrypted volumes, it automatically changes the encryption key. This feature improves security in Red Hat OpenStack Platform by not using the same encryption key more than once.

Image service manages deletion of encrypted keys

With this release, the Block Storage service (cinder) creates an encryption key in the Key Management service (barbican) when it uploads an encrypted volume to the Image service (glance). This creates a one-to-one relationship between an encryption key and a stored image. Encryption key deletion prevents unlimited resource consumption by the Key Management service.

Director supports back end availability zone configuration

This release adds director support to configure Block Storage back end availability zones. An availability zone is a provider-specific method of grouping cloud instances and services.

Removal of Data Processing service (sahara)

The Data Processing service (sahara) is deprecated in Red Hat OpenStack Platform (RHOSP) 15 and removed in RHOSP 16.0. Red Hat continues to offer support for the Data Processing service in RHOSP versions 13 and 15.

2.4. CEPH STORAGE

This section outlines the top new features for Ceph Storage.

Red Hat Ceph Storage Upgrade

To maintain compatibility with Red Hat Enterprise Linux 8, Red Hat OpenStack Platform 16.0 director deploys Red Hat Ceph Storage 4. You can use Red Hat OpenStack Platform 16.0 running on RHEL8 to connect to a preexisting external Red Hat Ceph Storage 3 cluster running on RHEL7.

2.5. CLOUD OPS

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

Service Telemetry Framework

Service Telemetry Framework (STF) provides the core components for a monitoring application framework for Red Hat OpenStack Platform. It is a data storage component deployed as an application on top of OpenShift 4.x and is managed by the Operator Lifecycle Manager. Data transport for metrics and events is provided using AMQ Interconnect. This first GA release of STF has the following features:

- Deployment of server-side Service Telemetry Framework as a micro-service application on OpenShift 4.3, leveraging the Operator Lifecycle Management (OLM)
- Full client-side installation capability through Red Hat OpenStack Platform director with collectd as the infrastructure data collector, Ceilometer for OpenStack events data, and AMQ Interconnect as the transport layer application
- Performance Metrics integration with Prometheus as the time-series database
- Events storage with ElasticSearch
- Alertmanager integration with a list of alerts out-of-the-box
- Infrastructure dashboard to visualize performance data with Grafana
- Multi-Cloud monitoring support with STF

2.6. TECHNOLOGY PREVIEWS

This section outlines features that are in technology preview in Red Hat OpenStack Platform 16.0.

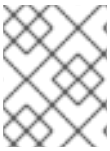

NOTE

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

2.6.1. New Technology Previews

Deploy and manage multiple overclouds from a single undercloud

This release includes the capability to deploy multiple overclouds from a single undercloud.

- Interact with a single undercloud to manage multiple distinct overclouds.
- Switch context on the undercloud to interact with different overclouds.
- Reduce redundant management nodes.

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.

Validation Framework

- Red Hat OpenStack Platform includes a validation framework to help verify the requirements and functionality of the undercloud and overcloud. The framework includes two types of validations:
- Manual Ansible-based validations, which you execute through the **openstack tripleo validator** command set.
- Automatic in-flight validations, which execute during the deployment process.
- Director provides a new set of commands to list validations and run validations against the undercloud and overcloud.
  These commands are:
  - **openstack tripleo validator list**
  - **openstack tripleo validator run**
    These commands interact with a set of Ansible-based tests from the **openstack-tripleo-validations** package. To enable this feature, set the **enable_validations** parameter to **true** in the **undercloud.conf** file and run **openstack undercloud install**.

New director feature to create an active-active configuration for Block Storage service

With Red Hat OpenStack Platform director, you can now deploy the Block Storage service (cinder) in an active-active configuration if the back end driver supports this configuration. As of GA, only the Ceph RADOS Block Device (RBD) back end driver supports an active-active configuration.
The new **cinder-volume-active-active.yaml** file defines the active-active cluster name by assigning a value to the **CinderVolumeCluster** parameter. **CinderVolumeCluster** is a global Block Storage

parameter, and prevents you from including clustered (active-active) and non-clustered back ends in the same deployment.

The **cinder-volume-active-active.yaml** file causes director to use the non-Pacemaker, cinder-volume Orchestration service template, and adds the etcd service to your Red Hat OpenStack Platform deployment as a distributed lock manager (DLM).

New director parameter for configuring Block Storage service availability zones

With Red Hat OpenStack Platform director, you can now configure different availability zones for Block Storage service (cinder) volume back ends. Director has a new parameter, **CinderXXXAvailabilityZone**, where XXX is associated with a specific back end.

New Redfish BIOS management interface for Bare Metal service

Red Hat OpenStack Platform Bare Metal service (ironic) now has a BIOS management interface, with which you can inspect and modify a device's BIOS configuration.
In Red Hat OpenStack Platform 16.0, the Bare Metal service supports BIOS management capabilities for data center devices that are Redfish API compliant. The Bare Metal service implements Redfish calls through the Python library, Sushy.

Deploying multiple Ceph clusters

You can use director to deploy multiple Ceph clusters, either on nodes dedicated to running Ceph or hyper-converged, using separate heat stacks for each cluster. For edge sites, you can deploy a hyper-converged infrastructure (HCI) stack that uses Compute and Ceph storage on the same node. For example, you might deploy two edge stacks named **HCI-01** and **HCI-02**, each in their own availability zone. As a result, each edge stack has its own Ceph cluster and Compute services.

New Compute (nova) configuration added to enable memoryBacking source type file, with shared access

A new Compute (nova) parameter is available, **QemuMemoryBackingDir**, which specifies the directory in which to store the memory backing file when a libvirt **memoryBacking** element is configured with **source type="file"** and **access mode="shared"**.
Note: The **memoryBacking** element is only available from libvirt 4.0.0 and QEMU 2.6.0.

Fencing for RedFish API

Fencing is now available with Pacemaker for RedFish API.

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.

Nova-less provisioning

In Red Hat OpenStack Platform 16.0, 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.

networking-ansible trunk port support

In Red Hat OpenStack Platform 16.0, you can use switch ports in trunk mode as well as access mode, and assign multiple VLANs to a switch port.

networking-ansible Arista support

In Red Hat OpenStack Platform 16.0, you can configure ML2 networking-ansible functionality with Arista Extensible Operating System (Arista EOS) switches.

Redfish virtual media boot

You can use Redfish virtual media boot to supply a boot image to the Baseboard Management Controller (BMC) of a node so that the BMC can insert the image into one of the virtual drives. The node can then boot from the virtual drive into the operating system that exists in the image.

CHAPTER 3. RELEASE INFORMATION

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. Notes for updates released during the support lifecycle of this Red Hat OpenStack Platform release will appear in the advisory text associated with each update.

3.1. RED HAT OPENSTACK PLATFORM 16.0 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#1716335

In Red Hat OpenStack Platform 16.0, live migrations with OVN enabled now succeed, as the flag, **live_migration_wait_for_vif_plug,** is enabled by default.
Previously, live migrations failed, because the system was waiting for OpenStack Networking (neutron) to send **vif_plugged** notifications.

BZ#1758302

Previously, the regular expression for the oslo.util library was not updated, and it failed to recognize the output format from a newer version of the emulator, qemu (version 4.1.0). This fix in Red Hat OpenStack 16.0 updates the regular expression, and the oslo.util.imageutils library now functions properly.

BZ#1769868

Previously, the mesh network infrastructure was configured incorrectly for the message router, QDR, and this caused the AMQP-1.0 message bus on the Service Telemetry Framework (STF) client not to function. This fix corrects the configuration for the qdrouterd daemon on all overcloud nodes, and the STF client now works properly.

BZ#1775246

The **NUMATopologyFilter** is now disabled when rebuilding instances. Previously, the filter would always execute and the rebuild would only succeed if a host had enough additional capacity for a second instance using the new image and existing flavor. This was incorrect and unnecessary behavior.

3.1.2. Enhancements

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

BZ#1222414

With this enhancement, support for live migration of instances with a NUMA topology has been added. Previously, this action was disabled by default. It could be enabled using the '[workarounds] enable_numa_live_migration' config option, but this defaulted to False because live migrating such instances resulted in them being moved to the destination host without updating any of the underlying NUMA guest-to-host mappings or the resource usage. With the new NUMA-aware live migration feature, if the instance cannot fit on the destination, the live migration will be attempted

on an alternate destination if the request is set up to have alternates. If the instance can fit on the destination, the NUMA guest-to-host mappings will be re-calculated to reflect its new host, and its resource usage updated.

BZ#1328124

Red Hat OpenStack Platform 16.0 director, now supports multi-compute cell deployments. With this enhancement, your cloud is better positioned for scaling out, because each individual cell has its own database and message queue on a cell controller and reduces the load on the central control plane. For more information, see "Scaling deployments with Compute cells" in the "Instances and Images" guide.

BZ#1360970

With this enhancement, support for live migration of instances with attached SR-IOV-based neutron interfaces has been added. Neutron SR-IOV interfaces can be grouped into two categories: direct mode and indirect mode. Direct mode SR-IOV interfaces are directly attached to the guest and exposed to the guest OS. Indirect mode SR-IOV interfaces have a software interface, for example, a macvtap, between the guest and the SR-IOV device. This feature enables transparent live migration for instances with indirect mode SR-IOV devices. As there is no generic way to copy hardware state during a live migration, direct mode migration is not transparent to the guest. For direct mode interfaces, mimic the workflow already in place for suspend and resume. For example, with SR-IOV devices, detach the direct mode interfaces before migration and re-attach them after the migration. As a result, instances with direct mode SR-IOV ports lose network connectivity during a migration unless a bond with a live migratable interface is created within the guest.
Previously, it was not possible to live migrate instances with SR-IOV-based network interfaces. This was problematic as live migration is frequently used for host maintenance and similar actions. Previously, the instance had to be cold migrated which involves downtime for the guest.

This enhancement results in the live migration of instances with SR-IOV-based network interfaces.

BZ#1463838

In Red Hat OpenStack Platform 16.0, it is now possible to specify QoS minimum bandwidth rules when creating network interfaces. This enhancement ensures that the instance is guaranteed a specified value of a network's available bandwidth. Currently, the only supported operations are resize and cold migrate.

BZ#1545700

The Red Hat OpenStack Platform Block Storage service (cinder) now automatically changes the encryption keys when cloning volumes. Note, that this feature currently does not support using Red Hat Ceph Storage as a cinder back end.

BZ#1545855

In Red Hat OpenStack Platform 16.0, you are now able to push, list, delete, and show (show metadata) images on the local registry.

- To push images from remote repository to the main repository:

```
$ sudo openstack tripleo container image push docker.io/library/centos
```

- To list the contents of the repository:

```
$ openstack tripleo container image list
```

- To delete images:

```
$ sudo openstack tripleo container image delete
```

- To show metadata for an image:

```
$ openstack tripleo container image show
```

BZ#1593057

With this enhancement, overcloud node deletion requires user confirmation before the action will be performed to reduce the likelihood that the action is performed unintentionally. The **openstack overcloud node delete <node>** command requires a Y/n confirmation before the action executes. You can bypass this by adding **--yes** to the command line.

BZ#1601926

Starting with this update, OSP deployments have full encryption between all the OVN services. All OVN clients (ovn-controller, neutron-server and ovn-metadata-agent) now connect to the OVSDB server using Mutual TLS encryption.

BZ#1625244

The Placement service has been extracted from the Compute (nova) service. It is now deployed and managed by the director, and runs as an additional container on the undercloud and on overcloud controller nodes.

BZ#1628541

In the Red Hat OpenStack Platform 16.0 dashboard (horizon), there is now a new form for changing a user's password. This form automatically appears when a user tries to sign on with an expired password.

BZ#1649264

The Red Hat OpenStack Platform Orchestration service (heat) now includes a new resource type, OS::Glance::WebImage, used for creating an Image service (glance) image from a URL using the Glance v2 API. This new resource type replaces an earlier one, OS::Glance::Image.

BZ#1653834

This enhancement adds the boolean parameter **NovaComputeEnableKsm**. The parameter enables the ksm and ksmtuned service on compute nodes. You can set **NovaComputeEnableKsm** for each Compute role. Default: **False**.

BZ#1666973

In Red Hat OpenStack Platform 16.0, you can now add custom Red Hat Ceph Storage configuration settings to any section of ceph.conf. Previously, custom settings were allowed only in the **[global]** section of ceph.conf.

BZ#1689816

In Red Hat OpenStack Platform 16.0, a new Orchestration service (heat) deployment parameter is available that enables administrators to turn on the nova metadata service on cell controllers:

```
parameter_defaults:
  NovaLocalMetadataPerCell: True
```

This new parameter automatically directs traffic from the OVN metadata agent on the cell computes to the nova metadata API service hosted on the cell controllers.

Depending on the RHOSP topology, the ability to run the metadata service on cell controllers can reduce the traffic on the central control plane.

BZ#1691025

You can now use the Octavia API to create a VIP access control list (ACL) to limit incoming traffic to a listener to a set of allowed source IP addresses (CIDRs). Any other incoming traffic is rejected. For more information, see "Secure a load balancer with an access control list" in the "Networking Guide."

BZ#1693372

With this enhancement, you can schedule dedicated (pinned) and shared (unpinned) instances on the same Compute node using the following parameters:

- **NovaComputeCpuDedicatedSet**- A comma-separated list or range of physical host CPU numbers to which processes for pinned instance CPUs can be scheduled. Replaces the NovaVcpuPinSet parameter, which is now deprecated.

- **NovaComputeCpuSharedSet** - A comma-separated list or range of physical host CPU numbers used to provide vCPU inventory, determine the host CPUs that unpinned instances can be scheduled to, and determine the host CPUS that instance emulator threads should be offloaded to for instances configured with the share emulator thread policy, **hw:emulator_threads_policy=share**. Note: This option previously existed but its purpose has been extended with this feature.
  It is no longer necessary to use host aggregates to ensure these instance types run on separate hosts. Also, the **[DEFAULT] reserved_host_cpus** config option is no longer necessary and can be unset.

  To upgrade:

  - For hosts that were previously used for pinned instances, the value of **NovaVcpuPinSet** should be migrated to **NovaComputeCpuDedicatedSet**.

  - For hosts that were previously used for unpinned instances, the value of **NovaVcpuPinSet** should be migrated to **NovaComputeCpuSharedSet**.

  - If there is no value set for **NovaVcpuPinSet**, then all host cores should be assigned to either **NovaComputeCpuDedicatedSet** or **NovaComputeCpuSharedSet**, depending on the type of instance running there.
    Once the upgrade is complete, it is possible to start setting both options on the same host. However, to do this, the host should be drained of instances as nova will not start when cores for an unpinned instance are not listed in **NovaComputeCpuSharedSet** and vice versa.

BZ#1696663

This update allows you to configure NUMA affinity for most neutron networks. This helps you ensure that instances are placed on the same host NUMA node as the NIC providing external connectivity to the vSwitch. You can configure NUMA affinity on networks that use: **--'provider:network_type'** of **'flat'** or **'vlan'** and a 'provider:physical_network' (L2 networks) or **--'provider:network_type'** of **'vxlan'**, **'gre' or 'geneve'** (L3 networks).

BZ#1700396

In Red Hat OpenStack Platform 16.0, you can now use director to specify an availability zone for the Block Storage service (cinder) back end type.

BZ#1767481

Previously, when Novajoin lost its connection to the IPA server, it would immediately attempt to reconnect. Consequently, timing issues could arise and prevent the connection from being re-established. With this update, you can use **retry_delay** to set the number of seconds to wait before retrying the IPA server connection. As a result, this is expected to help mitigate the timing issues.

BZ#1775575

You can now configure PCI NUMA affinity on an instance-level basis. This is required to configure NUMA affinity for instances with SR-IOV-based network interfaces. Previously, NUMA affinity was only configurable at a host-level basis for PCI passthrough devices.

BZ#1784806

In Red Hat Openstack Platform 16.0, a deployment enhancement eases configuring OVS-DPDK by automatically deriving the Orchestration service (heat) parameters required for the compute node on which OVS-DPDK is deployed. The Workflow service (mistral) has been enhanced to read heat templates and introspection data to automatically derive the necessary values for the heat parameters, NovaComputeCpuDedicatedSet and NovaComputeCpuSharedSet.

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#1228474

After a Red Hat OpenStack Platform 16.0 director deployment, the Identity service (keystone) now has a new default role, reader, which the other OpenStack services have not yet implemented. The reader role in keystone should not be used in a production environment, because the role is in technology preview and incorrectly grants privileges that users assigned to the role should not have, such as the ability to create volumes.

BZ#1288155

Defining multiple route tables and assigning routes to particular tables is a technology preview in Red Hat OpenStack Platform 16.0.
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 on a per-interface basis, as shown in this example:

```
network_config:
  -
    type: route_table
    name: custom
    table_id: 200
  -
    type: interface
    name: em1
    use_dhcp: false
    addresses:
    -
      ip_netmask: 192.0.2.1/24
    routes:
      -
        ip_netmask: 10.1.3.0/24
        next_hop: 192.0.2.5
        table: 200  # Use table ID or table name
    rules:
      - rule: "iif em1 table 200"
        comment: "Route incoming traffic to em1 with table 200"
      - rule: "from 192.0.2.0/24 table 200"
        comment: "Route all traffic from 192.0.2.0/24 with table 200"
      - rule: "add blackhole from 172.19.40.0/24 table 200"
      - rule: "add unreachable iif em1 from 192.168.1.0/24"
```

BZ#1375207

Previously, when using Red Hat Ceph Storage as a back end for both the Block Storage service (cinder) volumes and backups, any attempt to perform a full backup—after the first full backup—instead resulted in an incremental backup without any warning. In Red Hat OpenStack Platform 16.0, a technology preview has fixed this issue.

BZ#1459187

In Red Hat OpenStack Platform 16.0, a technology preview has been added to the Bare Metal Provisioning service (ironic) for deploying the overcloud on an IPv6 provisioning network. For more information, see "Configuring a custom IPv6 provisioning network," in the Bare Metal Provisioning guide.

BZ#1474394

In Red Hat OpenStack Platform 16.0, a technology preview has been added for the Bare Metal Provisioning service (ironic) deploying over an IPv6 provisioning network for BMaaS (Bare Metal as-a-Service) tenants.

BZ#1575079

In Red Hat OpenStack Platform 16.0, 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.

BZ#1593828

In Red Hat OpenStack Platform 16.0, a technology preview has been added for booting bare metal machines from virtual media using the Bare Metal Provisioning service (ironic).
If the baseboard management controller (BMC) for a machine supports Redfish hardware management protocol and virtual media service, ironic can instruct the BMC to pull a bootable image and "insert" it into a virtual drive on a node. The node can then boot from that virtual drive into the operating system residing on the image. Ironic hardware types based on the Redfish API support deploy, rescue (with a limitation), and boot (user) images over virtual media.

The major advantage of virtual media boot is that the insecure and unreliable TFTP image transfer phase of the PXE boot protocol suite is replaced by secure HTTP transport.

BZ#1600967

In Red Hat OpenStack Platform 16.0, a Workflow service (mistral) task is in technology preview that allows you to implement password rotation by doing the following:
Execute the rotate-password workflow to generate new passwords and store them in the plan environment.

Redeploy your overcloud.

You can also obtain your passwords after you have changed them.

To implement password rotation, follow these steps:

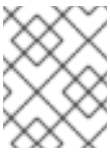

NOTE

The workflow task modifies the default passwords. The task does not modify passwords that are specified in a user-provided environment file.

1. Execute the new workflow task to regenerate the passwords:

```
$ source ./stackrc
$ openstack workflow execution create tripleo.plan_management.v1.rotate_passwords
'{"container": "overcloud"}'
```

This command generates new passwords for all passwords except for BarbicanSimpleCryptoKek and KeystoneFernet* and KeystoneCredential*. There are special procedures to rotate these passwords.

It is also possible to specify specific passwords to be rotated. The following command rotates only the specified passwords.

```
$ openstack workflow execution create tripleo.plan_management.v1.rotate_passwords
'{"container": "overcloud", "password_list": ["BarbicanPassword", "SaharaPassword",
"ManilaPassword"]}'
```

2. Redeploy your overcloud:

```
$ ./overcloud-deploy.sh
```

To retrieve the passwords, including the newly generated ones, follow these steps:

3. Run the following command:

```
$ openstack workflow execution create tripleo.plan_management.v1.get_passwords
'{"container": "overcloud"}'
```

You should see output from the command, similar to the following:

```
+--------------------+--------------------------------------------+
| Field              | Value                                      |
+--------------------+--------------------------------------------+
| ID                 | edcf9103-e1a8-42f9-85c1-e505c055e0ed       |
| Workflow ID        | 8aa2ac9b-22ee-4e7d-8240-877237ef0d0a       |
| Workflow name      | tripleo.plan_management.v1.rotate_passwords |
| Workflow namespace |                                            |
| Description        |                                            |
| Task Execution ID  | <none>                                     |
| Root Execution ID  | <none>                                     |
| State              | RUNNING                                    |
| State info         | None                                       |
| Created at         | 2020-01-22 15:47:57                        |
| Updated at         | 2020-01-22 15:47:57                        |
+--------------------+--------------------------------------------+
```

In the earlier example output, the value of State is RUNNING. State should eventually read SUCCESS.

4. Re-check the value of State:

```
$ openstack workflow execution show edcf9103-e1a8-42f9-85c1-e505c055e0ed
```

```
+--------------------+--------------------------------------------+
| Field              | Value                                      |
+--------------------+--------------------------------------------+
```

```
| ID                  | edcf9103-e1a8-42f9-85c1-e505c055e0ed        |
| Workflow ID         | 8aa2ac9b-22ee-4e7d-8240-877237ef0d0a        |
| Workflow name       | tripleo.plan_management.v1.rotate_passwords |
| Workflow namespace |                                              |
| Description         |                                             |
| Task Execution ID   | <none>                                      |
| Root Execution ID   | <none>                                      |
| State               | SUCCESS                                     |
| State info          | None                                        |
| Created at          | 2020-01-22 15:47:57                         |
| Updated at          | 2020-01-22 15:48:39                         |
+---------------------+---------------------------------------------+
```

5. When the value of State is SUCCESS, you can retrieve passwords:

```
$ openstack workflow execution output show edcf9103-e1a8-42f9-85c1-e505c055e0ed
```

You should see output similar to the following:

```
{
    "status": "SUCCESS",
    "message": {
        "AdminPassword": "FSn0sS1aAHp8YK2fU5niM3rxu",
        "AdminToken": "dTP0Wdy7DtblG80M54r4a2yoC",
        "AodhPassword": "fB5NQdRe37BaBVEWDHVuj4etk",
        "BarbicanPassword": "rn7yk7KPafKw2PWN71MvXpnBt",
        "BarbicanSimpleCryptoKek": "lrC3sGlV7-D7-
V_PI4vbDfF1Ujm5OjnAVFcnihOpbCg=",
        "CeilometerMeteringSecret": "DQ69HdlJobhnGWoBC0jM3drPF",
        "CeilometerPassword": "qI6xOpofuiXZnG95iUe8Oxv5d",
        "CephAdminKey": "AQDGVPpdAAAAABAAZMP56/VY+zCVcDT81+TOjg==",
        "CephClientKey": "AQDGVPpdAAAAABAAanYtA0ggpcoCbS1nLeDN7w==",
        "CephClusterFSID": "141a5ede-21b4-11ea-8132-52540031f76b",
        "CephDashboardAdminPassword":
"AQDGVPpdAAAAABAAKhsx630YKDhQrocS4o4KzA==",
        "CephGrafanaAdminPassword":
"AQDGVPpdAAAAABAAKBojG+CO72B0TdBRR0paEg==",
        "CephManilaClientKey": "AQDGVPpdAAAAABAAA1TVHrTVCC8xQ4skG4+d5A=="
    }
}
```

BZ#1621701

In Red Hat OpenStack Platform 16.0, a technology preview is added to the OpenStack Bare Metal service (ironic) to configure ML2 networking-ansible functionality with Arista Extensible Operating System (Arista EOS) switches. For more information, see "Enabling networking-ansible ML2 functionality," in the Bare Metal Provisioning guide.

BZ#1622233

In Red Hat OpenStack Platform 16.0, a technology preview has been added to modify switch ports to put them into trunking mode and assign more than one VLAN to them.

BZ#1623152

In Red Hat OpenStack Platform 16.0, a technology preview has been added to the Orchestration service (heat) for rsyslog changes:

- Rsyslog is configured to collect and forward container logs to be functionally equivalent to the fluentd installation.
- Administrators can configure rsyslog log forwarding in the same way as fluentd.

BZ#1628061

In Red Hat OpenStack Platform 16.0, you can use director to include in-flight validations in the service template. This feature is a technology preview in RHOSP 16.0. Additions can be inserted at the end of the step to be checked, or at the beginning of the next step.
In this example, a validation is performed to ensure that the rabbitmq service is running after its deployment:

```
deploy_steps_tasks:
  # rabbitmq container is supposed to be started during step 1
  # so we want to ensure it's running during step 2
  - name: validate rabbitmq state
    when: step|int == 2
    tags:
      - opendev-validation
      - opendev-validation-rabbitmq
    wait_for_connection:
      host: {get_param: [ServiceNetMap, RabbitmqNetwork]}
      port: 5672
      delay: 10
```

Heat enables you to include existing validations from the openstack-tripleo-validations roles:

```
deploy_steps_tasks:
  - name: some validation
    when: step|int == 2
    tags:
      - opendev-validation
      - opendev-validation-rabbitmq
    include_role:
      role: rabbitmq-limits
    # We can pass vars to included role, in this example
    # we override the default min_fd_limit value:
    vars:
      min_fd_limit: 32768
```

You can find the definition of the rabbitmq-limits role here: https://opendev.org/openstack/tripleo-validations/src/branch/stable/train/roles/rabbitmq_limits/tasks/main.yml

Here is an example of using the existing service health check:

```
deploy_steps_tasks:
  # rabbitmq container is supposed to be started during step 1
  # so we want to ensure it's running during step 2
  - name: validate rabbitmq state
    when: step|int == 2
    tags:
      - opendev-validation
```

```
  - opendev-validation-rabbitmq
command: >
  podman exec rabbitmq /openstack/healthcheck
```

BZ#1699449

Red Hat OpenStack Platform director now offers a technology preview for fence_redfish, a fencing agent for the Redfish API.

BZ#1700083

In Red Hat OpenStack Platform 16.0, a technology preview has been added for the Bare Metal Provisioning service (ironic) to work with Intel Speed Select processors.

BZ#1703956

In Red Hat OpenStack Platform 16.0, the Load-balancing service (octavia) now has a technology preview for UDP protocol.

BZ#1706896

In Red Hat OpenStack Platform 16.0, a technology preview has been added to the Image service (glance) that pre-caches images so that operators can warm the cache before they boot an instance.

BZ#1710089

Director has added the **overcloud undercloud minion install** command that you can use to configure an additional host to augment the Undercloud services.

BZ#1710092

Director now provides the ability to deploy an additional node that you can use to add additional heat-engine resources for deployment related actions.

BZ#1710093

Red Hat OpenStack Platform director now enables you to deploy an additional node that can be used to add additional Bare Metal Provisioning conductor service resources for system provisioning during deployments.

BZ#1710634

In Red Hat OpenStack Platform 16.0, a technology preview has been added to the Orchestration service (heat). A new parameter, **NovaSchedulerQueryImageType,** has been added that controls the Compute service (nova) placement and scheduler components query placement for image type (scheduler/query_placement_for_image_type_support).
When set to true (the default), **NovaSchedulerQueryImageType** excludes compute nodes that do not support the disk format of the image used in a boot request.

For example, the libvirt driver uses Red Hat Ceph Storage as an ephemeral back end, and does not support qcow2 images (without an expensive conversion step). In this case, enabling **NovaSchedulerQueryImageType** ensures that the scheduler does not send requests to boot a qcow2 image to compute nodes that use Red Hat Ceph Storage.

BZ#1749483

You can now forward the traffic from a TCP, UDP, or other protocol port of a floating IP address to a TCP, UDP, or other protocol port associated to one of the fixed IP addresses of a neutron port. Forwarded traffic is managed by an extension to the neutron API and by an OpenStack Networking plug-in. A floating IP address can have more than one forwarding definition configured. However, you cannot forward traffic for IP addresses that have a pre-existing association to an OpenStack Networking port. Traffic can only be forwarded for floating IP addresses that are managed by centralized routers on the network (legacy, HA, and DVR+HA).
To forward traffic for a port of a floating IP address, use the following OpenStack Networking plug-in command:

```
openstack floating ip port forwarding create
--internal-ip-address <internal-ip-address>
--port <port>
--internal-protocol-port <port-number>
--external-protocol-port <port-number>
--protocol <protocol>
<floating-ip>
```

--internal-ip-address <internal-ip-address> The fixed, IPv4, internal IP address of the neutron port that will receive the forwarded traffic.

--port <port> The name or ID of the neutron port that will receive the forwarded traffic.

--internal-protocol-port <port-number> The protocol port number of the neutron, fixed IP address that will receive the forwarded traffic.

--external-protocol-port <port-number> The protocol port number of the port of the floating IP address that will forward its traffic.

--protocol <protocol> The protocol that the port of the floating IP address uses (for example, TCP, UDP).

<floating-ip> The floating IP (IP address or ID) of the port that will forward its traffic.

Here is an example:

```
openstack floating ip port forwarding create \
  --internal-ip-address 192.168.1.2 \
  --port f7a08fe4-e79e-4b67-bbb8-a5002455a493 \
  --internal-protocol-port 18343 \
  --external-protocol-port 8343 \
  --protocol tcp \
  10.0.0.100
```

3.1.4. 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#1481814

Previously, when an encrypted Block Storage service (cinder) volume image was deleted, its corresponding key was not deleted.

In Red Hat OpenStack Platform 16.0, this issue has been resolved. When the Image service deletes a cinder volume image, it also deletes the key for the image.

BZ#1783044

With the general availability of Red Hat Ceph Storage version 4, you can now install **ceph-ansible** from the **rhceph-4-tools-for-rhel-8-x86_64-rpms** repository.

3.1.5. Known Issues

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

BZ#1574431

There is a known issue for the Block Storage service (cinder) where quota commands do not work as expected. The cinder CLI allows users to successfully create quota entries without checking 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, CLI users should make sure to specify a valid project ID when creating quota entries, and monitor cinder for dummy records.

BZ#1647005

Nova-compute ironic driver tries to update BM node while the node is being cleaned up. The cleaning takes approximately five minutes but nova-compute attempts to update the node for approximately two minutes. After timeout, nova-compute stops and puts nova instance into ERROR state.
As a workaround, set the following configuration option for nova-compute service:

```
[ironic]
api_max_retries = 180
```

As a result, nova-compute continues to attempt to update BM node longer and eventually succeeds.

BZ#1734301

Currently, the OVN load balancer does not open new connections when fetching data from members. The load balancer modifies the destination address and destination port and sends request packets to the members. As a result, it is not possible to define an IPv6 member while using an IPv4 load balancer address and vice versa. There is currently no workaround for this issue.

BZ#1769880

There is a known issue where migrations from ML2/OVS to OVN fail. The failure is caused by the new protective mechanism in Red Hat OpenStack Platform director to prevent upgrades while changing mechanism drivers.
For the workaround, see "Preparing for the migration," in the "Networking with Open Virtual Network" guide.

BZ#1779221

Red Hat OpenStack Platform deployments that use the Linux bridge ML2 driver and agent are unprotected against Address Resolution Protocol (ARP) spoofing. The version of Ethernet bridge frame table administration (ebtables) that is part of Red Hat Enterprise Linux 8 is incompatible with the Linux bridge ML2 driver.
The Linux Bridge ML2 driver and agent were deprecated in Red Hat OpenStack Platform 11, and should not be used.

Red Hat recommends that you use instead the ML2 Open Virtual Network (OVN) driver and services, the default deployed by Red Hat OpenStack Platform director.

BZ#1789822

Replacement of an overcloud Controller might cause swift rings to become inconsistent across nodes. This can result in decreased availability of Object Storage service. If this happens, 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#1790467

There is a known issue where Red Hat OpenStack Platform 16.0, where metadata information required for configuring OpenStack instances is not available, and they might be started without connectivity.
An ordering issue causes the haproxy wrappers not to be updated for the ovn_metadata_agent.

A possible workaround is for the cloud operator to run the following Ansible command to restart the ovn_metadata_agent on select nodes after the update, to ensure that the ovn_metadata_agent is using an updated version of the haproxy wrapper script:

```
`ansible -b <nodes> -i /usr/bin/tripleo-ansible-inventory -m shell -a "status=`sudo systemctl is-
active tripleo_ovn_metadata_agent`; if test \"$status\" == \"active\"; then sudo systemctl restart
tripleo_ovn_metadata_agent; echo restarted; fi"`
```

In the earlier Ansible command, **nodes** may be a single node (for example, **compute-0**), all computes (for example, **compute***) or **"all"**.

As the ovn_metadata_agent is most commonly found on compute nodes, the following Ansible command restarts the agent for all compute nodes in the cloud:

```
`ansible -b compute* -i /usr/bin/tripleo-ansible-inventory -m shell -a "status=`sudo systemctl is-
active tripleo_ovn_metadata_agent`; if test \"$status\" == \"active\"; then sudo systemctl restart
tripleo_ovn_metadata_agent; echo restarted; fi"`
```

After you restart the ovn_metadata_agent services, they use the updated haproxy wrapper script, which enables them to provide metadata to VMs when they are started. Affected VMs already running should behave normally when they are restarted after the workaround has been applied.

BZ#1793166

There is a known issue in Red Hat OpenStack 16.0, where KVM guests do not start on IBM POWER8 systems unless the simultaneous multithreading (SMT) control is disabled. SMT is not disabled automatically.
The workaround is to execute **sudo ppc64_cpu --smt=off** on any IBM POWER8 compute nodes after deploying the overcloud, and any subsequent reboots.

BZ#1793440

In Red Hat OpenStack 16.0, there is a known issue where the command, "openstack network agent list," intermittently indicates that the OVN agents are down, when the agents are actually alive and the cloud is operational.
The affected agents are: OVN Controller agent, OVN Metadata agent, and OVN Controller Gateway agent.

There is currently no workaround for this issue. You should ignore the output of the "openstack network agent list" command.

BZ#1794328

There is a known issue where Red Hat OpenStack Platform 16.0 overcloud installs fail, when the Load-balancing service (octavia) is configured with a composable role. Currently, there is no identified workaround for this issue. For more information, see the BZ# itself: https://bugzilla.redhat.com/show_bug.cgi?id=1794328.

BZ#1795165

There is a known issue for OpenStack Networking (neutron) where all instances created inherit the dns_domain value associated with the network, and not the dns_domain value configured for an internal DNS.
The cause of this issue is that the network dns_domain attribute overrides the neutron dns_domain config option.

To avoid this issue, do not set the dns_domain attribute for the network, if you want to use the internal DNS feature.

BZ#1795688

To allow Placement services deployed on the Controller node to be accessed by the neutron_api service, as required when using the Novacontrol role, add the following hieradata configuration to your Controller environment file:

```
service_config_settings:
    placement:
        neutron::server::placement::password: <Nova password>
  neutron::server::placement::www_authenticate_uri: <Keystone Internal API URL>
  neutron::server::placement::project_domain_name: 'Default'
  neutron::server::placement::project_name: 'service'
  neutron::server::placement::user_domain_name: 'Default'
  neutron::server::placement::username: nova
  neutron::server::placement::auth_url: <Keystone Internal API URL>
  neutron::server::placement::auth_type: 'password'
  neutron::server::placement::region_name: <Keystone Region>
```

For more information on using Puppet to customizing hieradata for roles, see https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.0/html-single/advanced_overcloud_customization/index#sect-Customizing_Puppet_Configuration_Data.

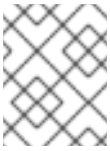

NOTE

This configuration is required when deploying an overcloud with a custom role where Placement is not running on the same nodes as neutron_api.

BZ#1795956

There is a known issue for the Red Hat OpenStack Platform Load-balancing service: the containers octavia_api and octavia_driver_agent fail to start when rebooting a node.
The cause for this issue is that the directory, /var/run/octavia, does not exist when the node is rebooted.

To fix this issue, add the following line to the file, /etc/tmpfiles.d/var-run-octavia.conf:

```
d /var/run/octavia 0755 root root - -
```

BZ#1796215

In Red Hat OpenStack Platform 16.0, there is a known issue when ansible-playbook can sometimes fail during configuration of the overcloud nodes. The cause for the failure is the tripleo-admin user is not authorized for ssh. Furthermore, an **openstack overcloud deploy** command argument, **--stack-only**, no longer runs the enable ssh admin workflow to authorize the tripleo-admin user.
The workaround is to use the **openstack overcloud admin authorize** command to run the enable ssh admin workflow on its own when using **--stack-only** and the manual config-download commands. for more information, see "Separating the provisioning and configuration processes" in the Director Installation and Usage guide.

BZ#1797047

The manila access-list feature requires Red Hat Ceph Storage 4.1 or later. Red Hat Ceph Storage 4.0 has a packaging issue. As a result, customers cannot use manila access-list. Share creation works, but without manila access-list, the share is unusable. Consequently, customers cannot use the Shared File System service with CephFS via NFS. For more information,
see https://bugzilla.redhat.com/show_bug.cgi?id=1797075.

BZ#1797892

There is a known issue in Red Hat OpenStack Platform 16.0, when nodes experiencing hard (ungraceful) shutdowns put containers—that were previously running---in a "Created" state in podman when the node is turned back on.
The reason for this issue is that the metadata agent fails to spawn a new container because it already exists in the "Created" state. The haproxy side-car container wrapper script expects containers to be in only the "Exited" state, and does not cleanup containers in the "Created" state.

The possible workaround is for the cloud operator to run the following Ansible ad-hoc command to clean up all haproxy containers in the "Created" state. You must run this Ansible ad-hoc command from the undercloud on particular node, on a group of nodes, or on the whole cluster:

```
`ansible -b <nodes> -i /usr/bin/tripleo-ansible-inventory -m shell -a "podman ps -a --format {{'{{'}}.ID{{'}}'}} -f name=haproxy,status=created | xargs podman rm -f || :"`
```

In the earlier Ansible ad-hoc command, **nodes** can be a single host from the inventory, a group of hosts, or "all".

Here is an example of running the command on **compute-0**:

```
`ansible -b compute-0 -i /usr/bin/tripleo-ansible-inventory -m shell -a "podman ps -a --format {{'{{'}}.ID{{'}}'}} -f name=haproxy,status=created | xargs podman rm -f || :"`
```

After running the Ansible ad-hoc command, the metadata-agent should then spawn a new container for the given network.

3.1.6. Removed Functionality

BZ#1518222

In Red Hat OpenStack Platform 16.0, a part of the Telemetry service, the ceilometer client (that was deprecated in an earlier RHOSP release) is no longer supported and has been removed. Note that ceilometer continues to be a part of RHOSP as an agent-only service (no client and no API).

BZ#1631508

In Red Hat OpenStack Platform 16.0, the controller-v6.yaml file is removed. The routes that were defined in controller-v6.yaml are now defined in controller.yaml. (The controller.yaml file is a NIC configuration file that is rendered from values set in roles_data.yaml.)
Previous versions of Red Hat OpenStack director, included two routes: one for IPv6 on the External network (default) and one for IPv4 on the Control Plane.

To use both default routes, make sure that the controller definition in roles_data.yaml contains both networks in default_route_networks (for example, **default_route_networks: ['External', 'ControlPlane']**).

BZ#1712981

The Data Processing service (sahara) is deprecated in Red Hat OpenStack Platform (RHOSP) 15 and removed in RHOSP 16.0. Red Hat continues to offer support for the Data Processing service in RHOSP versions 13 and 15.

BZ#1754560

In Red Hat OpenStack Platform 16.0, the Elastic Compute Cloud (EC2) API is no longer supported. The EC2 API support is now deprecated in director and will be removed in a future RHOSP release.

BZ#1764894

In Red Hat OpenStack Platform 16.0, the following environment file has been removed: **/usr/share/openstack-tripleo-heat-templates/environments/deployed-server-bootstrap-environment-rhel.yaml**.
This environment file was previously used when using pre-provisioned nodes. It was deprecated in a previous RHOSP release, and now it has been removed.

3.2. RED HAT OPENSTACK PLATFORM 16.0.1 MAINTENANCE RELEASE 4 MARCH 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. Enhancements

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

BZ#1784222

With this update, the pcs service now restricts listening to the InternalApi network by default.

BZ#1790752

Previously, when using Red Hat Ceph Storage as a back end for both the Block Storage service (cinder) volumes and backups, any attempt to perform a full backup—after the first full backup—instead resulted in an incremental backup without any warning. In Red Hat OpenStack Platform 16.0.1, the fix for this issue is fully supported.

3.2.2. Known Issues

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

BZ#1769880

There is a known issue where migrations from ML2/OVS to OVN fail. The failure is caused by the new protective mechanism in Red Hat OpenStack Platform director to prevent upgrades while changing mechanism drivers.

For the workaround, see "Preparing for the migration" in the Networking with Open Virtual Network guide: https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.0/html/networking_with_open_virtual_network/migrating-ml2ovs-to-ovn#preparing_for_the_migration

BZ#1790467

There is a known issue in Red Hat OpenStack Platform 16.0, where metadata information required for configuring OpenStack instances is not available, and instances might be started without connectivity.

An ordering issue causes the haproxy wrappers not to be updated for the ovn_metadata_agent service.

Workaround: Run the following Ansible command to restart the ovn_metadata_agent service on select nodes after the update to ensure that the ovn_metadata_agent service uses an updated version of the haproxy wrapper script:

**ansible -b <nodes> -i /usr/bin/tripleo-ansible-inventory -m shell -a "status=`sudo systemctl is-active tripleo_ovn_metadata_agent**; if test \"$status\" == \"active\"; then sudo systemctl restart tripleo_ovn_metadata_agent; echo restarted; fi"`

In this command, **nodes** can be a single node (for example, **compute-0**), all Compute nodes (for example, **compute***) or **"all"**.

After you restart the ovn_metadata_agent services, the services use the updated haproxy wrapper script and can provide metadata to VMs at startup. After you apply the workaround, affected VMs that are already running behave normally after a restart.

BZ#1793440

In Red Hat OpenStack 16.0, there is a known issue where the command **openstack network agent list** intermittently indicates that the OVN agents are down, when the agents are actually alive and the cloud is operational.

The affected agents are: OVN Controller agent, OVN Metadata agent, and OVN Controller Gateway agent.

There is currently no workaround for this issue. Ignore the output of the "openstack network agent list" command.

BZ#1795165

There is a known issue for OpenStack Networking (neutron) where all instances created inherit the dns_domain value associated with the network, and not the dns_domain value configured for an internal DNS.

The cause of this issue is that the network dns_domain attribute overrides the neutron dns_domain config option.

To avoid this issue, do not set the dns_domain attribute for the network if you want to use the internal DNS feature.

BZ#1795688

To allow the neutron_api service to access Placement services on Controller nodes, for example, when you use the Novacontrol role, add the following hieradata configuration to your Controller environment file:

```
service_config_settings:
    placement:
        neutron::server::placement::password: <Nova password>
  neutron::server::placement::www_authenticate_uri: <Keystone Internal API URL>
  neutron::server::placement::project_domain_name: 'Default'
  neutron::server::placement::project_name: 'service'
  neutron::server::placement::user_domain_name: 'Default'
  neutron::server::placement::username: nova
  neutron::server::placement::auth_url: <Keystone Internal API URL>
  neutron::server::placement::auth_type: 'password'
  neutron::server::placement::region_name: <Keystone Region>
```

For more information about using Puppet to customize hieradata for roles, see https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.0/html-single/advanced_overcloud_customization/index#sect-Customizing_Puppet_Configuration_Data.

Note: This configuration is required when deploying an overcloud with a custom role where Placement is not running on the same nodes as neutron_api.

BZ#1797892

There is a known issue in Red Hat OpenStack Platform 16.0, when nodes that experience a hard shutdown put containers that were previously running into a **Created** state in podman when the node reboots.

As a workaround, you can run the following Ansible command to clean all haproxy containers in the **Created** state:

ansible -b <nodes> -i /usr/bin/tripleo-ansible-inventory -m shell -a "podman ps -a --format {{'{{'}}.ID{{'}}'}} -f name=haproxy,status=created | xargs podman rm -f || :"

Replace **<nodes>** with a single host from the inventory, a group of hosts, or **all**. After you run this command, the metadata-agent spawns a new container for the given network.

BZ#1802573

There is a known issue where Mistral containers do not restart during minor updates and the update prepare times out after 10 hours.

The workaround is to restart the containers manually.

BZ#1804848

There is a known issue when all of the following conditions exist:

(0) You are using the OpenStack Train release (or code from master (Ussuri development))

(1) cinder_encryption_key_id and cinder_encryption_key_deletion_policy are not included in the non_inheritable_image_properties setting in nova.conf. These properties are not included by default.

(2) A user has created a volume of an encrypted volume-type in the Block Storage service (cinder). For example, Volume-1.

(3) Using the Block Storage service, the user has uploaded the encrypted volume as an image to the Image service (glance). For example, Image-1.

(4) Using the Compute service (nova), the user has attempted to boot a server from the image directly. Note: this is an unsupported action, the supported workflow is to use the image to boot-from-volume.

(5) Although an unsupported action, if a user does (4), it currently results in a server in status ACTIVE but which is unusable because the operating system cannot be found.

(6) Using the Compute service, the user requests the createImage action on the unusable server, resulting in the creation of Image-2.

(7) Using the Image service, the user deletes Image-2 which has inherited the cinder_encryption_key_* properties from Image-1 and the encryption key is deleted.

As a result, Image-1 is rendered non-decryptable so that it can no longer be used in the normal boot-from-volume workflow.

The workaround for this issue is to add the cinder_encryption_key_id,cinder_encryption_key_deletion_policy properties to the non_inheritable_image_properties option in the [DEFAULT] section of nova.conf. Image-2 can be deleted and the encryption key used by Image-1 remains available.

3.3. RED HAT OPENSTACK PLATFORM 16.0.2 MAINTENANCE RELEASE 14 MAY 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. Enhancements

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

BZ#1653834

This enhancement adds the Boolean parameter **NovaComputeEnableKsm**. The parameter enables the ksm and ksmtuned service on compute nodes. You can set **NovaComputeEnableKsm** for each Compute role. The default value is `False`.

BZ#1695898

Director operations involving the RADOS gateway no longer require interaction with puppet-ceph. Previously, tripleo-heat-templates had a dependency on puppet-ceph for the RADOS gateway component deployment. The move to tripleo-ansible eliminates this dependency.

BZ#1696717

This feature enables the Red Hat OpenStack Platform director to deploy the Shared File System (manila) with an external Ceph Storage cluster. In this type of deployment, Ganesha still runs on the Controller nodes that Pacemaker manages using an active-passive configuration.
This feature is supported with Ceph Storage 4.1 or later.

BZ#1749483

In the second maintenance release of Red Hat OpenStack Platform 16.0, IP port forwarding for OVS/ML2 has moved from technical preview to being fully supported. For more information, see the floating ip port forwarding create command in the *Command Line Interface Reference*.

BZ#1777052

The Service Telemetry Framework (STF) release v1.0 is now available for general availability. STF provides the core components for a monitoring application framework for Red Hat OpenStack Platform (RHOSP). It is a data storage component deployed as an application on top of OpenShift 4.x and is managed by the Operator Lifecycle Manager. Data transport for metrics and events is provided using AMQ Interconnect.
The release of STF v1.0 replaces and deprecates the Technology Preview version.

BZ#1790753

This update makes it possible for the Block Storage service (cinder) to attach Ceph RADOS block device (RBD) volumes to multiple instances simultaneously.

BZ#1790754

With this update, you can now enable Red Hat Ceph Storage Dashboard with the Red Hat OpenStack Platform director. The Red Hat Storage Ceph Dashboard is a built-in, web-based Ceph management and monitoring application to visualise and monitor various aspects in your cluster. Ceph Dashboard requires Red Hat Ceph Storage 4.1 or later.

BZ#1798917

A new Red Hat OpenStack Platform Orchestration service (heat) parameter controls whether the Block Storage service (cinder) flattens RADOS block device (RBD) volumes created from snapshots. Flattening a volume removes its dependency on the snapshot. If you set the value of **CinderRbdFlattenVolumeFromSnapshot** to true, cinder flattens RBD volumes. The default value of **CinderRbdFlattenVolumeFromSnapshot** and the cinder RBD driver is **false**.

3.3.2. 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#1703956

In Red Hat OpenStack Platform 16.0, the Load-balancing service (octavia) now has a technology preview for UDP protocol.

3.3.3. 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#1823835

RHOSP 16.0 works only with RHEL 8.1. Ensure that all the hosts of your OSP deployment are pinned to RHEL 8.1 before running the update.
See "Locking the environment to a Red Hat Enterprise Linux release" [1] in the guide "Keeping Red Hat OpenStack Platform Updated."

[1] https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.0/html/keeping_red_hat_openstack_platform_updated/preparing-for-a-minor-update#locking-the-environment-to-a-red-hat-enterprise-linux-release

3.3.4. Known Issues

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

BZ#1795956

There is a known issue for the Red Hat OpenStack Platform Load-balancing service: the containers octavia_api and octavia_driver_agent fail to start when rebooting a node.
The cause for this issue is that the directory, /var/run/octavia, does not exist when the node is rebooted.

To fix this issue, add the following line to the file, /etc/tmpfiles.d/var-run-octavia.conf:

d /var/run/octavia 0755 root root - -

BZ#1824093

A Grafana Ceph 4.1 dependency causes Ceph dashboard bugs. The Ceph dashboard requires Ceph 4.1 and a Grafana container based on ceph4-rhel8. Presently, Red Hat supports ceph3-rhel7.3. This discrepancy causes the following dashboard bugs:

- When you navigate to **Pools** > **Overall Performance**, Grafana returns the following error:

  > TypeError: l.c[t.type] is undefined
  > true

- When you view a pool's performance details (**Pools** > select a pool from the list > **Performance Details**) the Grafana bar is displayed along with other graphs and values, but it should not be there.
  These bugs will be fixed after rebasing to a newer Grafana version.

BZ#1837558

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

If you have ML2/OVN in an existing deployment, perform the following steps:

1. Set the **enable_distributed_floating_ip parameter** in the [ovs] section of neutron.conf to False. You should also set **NeutronEnableDVR: false** in an environment file used in any re-deployments so that the next re-deployment does not re-enable DVR.

2. Update the floating IP that requires external SNAT to work through the Neutron API (for example, by changing its description).

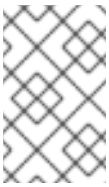

NOTE

Disabling DVR causes traffic to be centralized. All L3 traffic goes through the controller/network nodes. This may affect scale, data plane performance, and throughput.

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-2020:0283 – RED HAT OPENSTACK PLATFORM 16.0 GENERAL AVAILABILITY ADVISORY

The bugs contained in this section are addressed by advisory RHEA-2020:0283. Further information about this advisory is available at link: https://access.redhat.com/errata/RHEA-2020:0283.html.

Changes to the distribution component:

- In Red Hat OpenStack Platform 16.0, a part of the Telemetry service, the ceilometer client (that was deprecated in an earlier RHOSP release) is no longer supported and has been removed. Note that ceilometer continues to be a part of RHOSP as an agent-only service (no client and no API). (BZ#1518222)

Changes to the openstack-cinder component:

- Previously, when using Red Hat Ceph Storage as a back end for both the Block Storage service (cinder) volumes and backups, any attempt to perform a full backup—after the first full backup—instead resulted in an incremental backup without any warning. In Red Hat OpenStack Platform 16.0, a technology preview has fixed this issue. (BZ#1375207)

- The Red Hat OpenStack Platform Block Storage service (cinder) now automatically changes the encryption keys when cloning volumes. Note, that this feature currently does not support using Red Hat Ceph Storage as a cinder back end. (BZ#1545700)

Changes to the openstack-glance component:

- Previously, when an encrypted Block Storage service (cinder) volume image was deleted, its corresponding key was not deleted.
  In Red Hat OpenStack Platform 16.0, this issue has been resolved. When the Image service deletes a cinder volume image, it also deletes the key for the image. (BZ#1481814)

- In Red Hat OpenStack Platform 16.0, a technology preview has been added to the Image service (glance) that pre-caches images so that operators can warm the cache before they boot an instance. (BZ#1706896)

Changes to the openstack-heat component:

- The Red Hat OpenStack Platform Orchestration service (heat) now includes a new resource type, OS::Glance::WebImage, used for creating an Image service (glance) image from a URL using the Glance v2 API. This new resource type replaces an earlier one, OS::Glance::Image. (BZ#1649264)

Changes to the openstack-keystone component:

- Keystone now supports a basic set of default roles (for example, admin, member, and reader) that are present in the system after a Red Hat OpenStack Platform director deployment. These default roles are incorporated into the default director policies across all authorization targets (for example, system, domains, and projects). (BZ#1228474)

Changes to the openstack-manila component:

- In Red Hat OpenStack Platform 16.0, a technology preview has been added for the Shared File Systems service (manila) for IPv6 to work in the CephFS NFS driver. (BZ#1575079)

Changes to the openstack-neutron component:

- Red Hat OpenStack Platform deployments that use the Linux bridge ML2 driver and agent are unprotected against Address Resolution Protocol (ARP) spoofing. The version of Ethernet bridge frame table administration (ebtables) that is part of Red Hat Enterprise Linux 8 is incompatible with the Linux bridge ML2 driver.
  The Linux Bridge ML2 driver and agent were deprecated in Red Hat OpenStack Platform 11, and should not be used.

  Red Hat recommends that you use instead the ML2 Open Virtual Network (OVN) driver and services, the default deployed by Red Hat OpenStack Platform director. (BZ#1779221)

- You can now forward the traffic from a TCP, UDP, or other protocol port of a floating IP address to a TCP, UDP, or other protocol port associated to one of the fixed IP addresses of a neutron port. Forwarded traffic is managed by an extension to the neutron API and by an OpenStack Networking plug-in. A floating IP address can have more than one forwarding definition configured. However, you cannot forward traffic for IP addresses that have a pre-existing association to an OpenStack Networking port. Traffic can only be forwarded for floating IP addresses that are managed by centralized routers on the network (legacy, HA, and DVR+HA). To forward traffic for a port of a floating IP address, use the following OpenStack Networking plug-in command:

  ```
  openstack floating ip port forwarding create
  --internal-ip-address <internal-ip-address>
  --port <port>
  --internal-protocol-port <port-number>
  --external-protocol-port <port-number>
  --protocol <protocol>
  <floating-ip>
  ```

  --internal-ip-address <internal-ip-address> The fixed, IPv4, internal IP address of the neutron port that will receive the forwarded traffic.

  --port <port> The name or ID of the neutron port that will receive the forwarded traffic.

  --internal-protocol-port <port-number> The protocol port number of the neutron, fixed IP address that will receive the forwarded traffic.

  --external-protocol-port <port-number> The protocol port number of the port of the floating IP address that will forward its traffic.

  --protocol <protocol> The protocol that the port of the floating IP address uses (for example, TCP, UDP).

  <floating-ip> The floating IP (IP address or ID) of the port that will forward its traffic.

  Here is an example:

  openstack floating ip port forwarding create \ --internal-ip-address 192.168.1.2 \ --port f7a08fe4-e79e-4b67-bbb8-a5002455a493 \ --internal-protocol-port 18343 \ --external-protocol-port 8343 \ --protocol tcp \ 10.0.0.100 (BZ#1749483)

Changes to the openstack-nova component:

- With this enhancement, support for live migration of instances with a NUMA topology has been added. Previously, this action was disabled by default. It could be enabled using the '[workarounds] enable_numa_live_migration' config option, but this defaulted to False because live migrating such instances resulted in them being moved to the destination host without updating any of the underlying NUMA guest-to-host mappings or the resource usage. With the new NUMA-aware live migration feature, if the instance cannot fit on the destination, the live migration will be attempted on an alternate destination if the request is set up to have alternates. If the instance can fit on the destination, the NUMA guest-to-host mappings will be re-calculated to reflect its new host, and its resource usage updated. (BZ#1222414)

- With this enhancement, support for live migration of instances with attached SR-IOV-based neutron interfaces has been added. Neutron SR-IOV interfaces can be grouped into two categories: direct mode and indirect mode. Direct mode SR-IOV interfaces are directly attached to the guest and exposed to the guest OS. Indirect mode SR-IOV interfaces have a software interface, for example, a macvtap, between the guest and the SR-IOV device. This feature enables transparent live migration for instances with indirect mode SR-IOV devices. As there is no generic way to copy hardware state during a live migration, direct mode migration is not transparent to the guest. For direct mode interfaces, mimic the workflow already in place for suspend and resume. For example, with SR-IOV devices, detach the direct mode interfaces before migration and re-attach them after the migration. As a result, instances with direct mode SR-IOV ports lose network connectivity during a migration unless a bond with a live migratable interface is created within the guest.
  Previously, it was not possible to live migrate instances with SR-IOV-based network interfaces. This was problematic as live migration is frequently used for host maintenance and similar actions. Previously, the instance had to be cold migrated which involves downtime for the guest.

  This enhancement results in the live migration of instances with SR-IOV-based network interfaces. (BZ#1360970)

- In Red Hat OpenStack Platform 16.0, it is now possible to specify QoS minimum bandwidth rules when creating network interfaces. This enhancement ensures that the instance is guaranteed a specified value of a network's available bandwidth. Currently, the only supported operations are resize and cold migrate. (BZ#1463838)

- The **NUMATopologyFilter** is now disabled when rebuilding instances. Previously, the filter would always execute and the rebuild would only succeed if a host had enough additional capacity for a second instance using the new image and existing flavor. This was incorrect and unnecessary behavior. (BZ#1775246)

- You can now configure PCI NUMA affinity on an instance-level basis. This is required to configure NUMA affinity for instances with SR-IOV-based network interfaces. Previously, NUMA affinity was only configurable at a host-level basis for PCI passthrough devices. (BZ#1775575)

- With this enhancement, you can schedule dedicated (pinned) and shared (unpinned) instances on the same Compute node using the following parameters:

  - **NovaComputeCpuDedicatedSet**- A comma-separated list or range of physical host CPU numbers to which processes for pinned instance CPUs can be scheduled. Replaces the NovaVcpuPinSet parameter, which is now deprecated.

  - **NovaComputeCpuSharedSet** - A comma-separated list or range of physical host CPU numbers used to provide vCPU inventory, determine the host CPUs that unpinned instances can be scheduled to, and determine the host CPUS that instance emulator threads should be offloaded to for instances configured with the share emulator thread policy, **hw:emulator_threads_policy=share**. Note: This option previously existed but its purpose has been extended with this feature.

It is no longer necessary to use host aggregates to ensure these instance types run on separate hosts. Also, the **[DEFAULT] reserved_host_cpus** config option is no longer necessary and can be unset.

To upgrade:

- For hosts that were previously used for pinned instances, the value of **NovaVcpuPinSet** should be migrated to **NovaComputeCpuDedicatedSet**.
- For hosts that were previously used for unpinned instances, the value of **NovaVcpuPinSet** should be migrated to **NovaComputeCpuSharedSet**.
- If there is no value set for **NovaVcpuPinSet**, then all host cores should be assigned to either **NovaComputeCpuDedicatedSet** or **NovaComputeCpuSharedSet**, depending on the type of instance running there.
  Once the upgrade is complete, it is possible to start setting both options on the same host. However, to do this, the host should be drained of instances as nova will not start when cores for an unpinned instance are not listed in **NovaComputeCpuSharedSet** and vice versa. (BZ#1693372)

Changes to the openstack-octavia component:

- You can now use the Octavia API to create a VIP access control list (ACL) to limit incoming traffic to a listener to a set of allowed source IP addresses (CIDRs). Any other incoming traffic is rejected. For more information, see "Secure a load balancer with an access control list" in the "Networking Guide." (BZ#1691025)
- In Red Hat OpenStack Platform 16.0, the Load-balancing service (octavia) now has a technology preview for UDP protocol. (BZ#1703956)

Changes to the openstack-placement component:

- The Placement service has been extracted from the Compute (nova) service. It is now deployed and managed by the director, and runs as an additional container on the undercloud and on overcloud controller nodes. (BZ#1625244)

Changes to the openstack-tripleo-common component:

- In Red Hat OpenStack Platform 16.0, a Workflow service (mistral) task is in technology preview that allows you to implement password rotation by doing the following:
  - Execute the rotate-password workflow to generate new passwords and store them in the plan environment.
  - Redeploy your overcloud.
    You can also obtain your passwords after you have changed them.

  To implement password rotation, follow these steps:

  

  NOTE

  The workflow task modifies the default passwords. The task does not modify passwords that are specified in a user-provided environment file.

  1. Execute the new workflow task to regenerate the passwords:

```
$ source ./stackrc
$ openstack workflow execution create
tripleo.plan_management.v1.rotate_passwords '{"container": "overcloud"}'
```

This command generates new passwords for all passwords except for BarbicanSimpleCryptoKek and KeystoneFernet* and KeystoneCredential*. There are special procedures to rotate these passwords.

It is also possible to specify specific passwords to be rotated. The following command rotates only the specified passwords.

```
$ openstack workflow execution create
tripleo.plan_management.v1.rotate_passwords '{"container": "overcloud",
"password_list": ["BarbicanPassword", "SaharaPassword", "ManilaPassword"]}'
```

2. Redeploy your overcloud:

```
$ ./overcloud-deploy.sh
```

To retrieve the passwords, including the newly generated ones, follow these steps:

3. Run the following command:

```
$ openstack workflow execution create tripleo.plan_management.v1.get_passwords
'{"container": "overcloud"}'
```

You should see output from the command, similar to the following:

```
+--------------------+--------------------------------------------+
| Field              | Value                                      |
+--------------------+--------------------------------------------+
| ID                 | edcf9103-e1a8-42f9-85c1-e505c055e0ed       |
| Workflow ID        | 8aa2ac9b-22ee-4e7d-8240-877237ef0d0a       |
| Workflow name      | tripleo.plan_management.v1.rotate_passwords |
| Workflow namespace |                                            |
| Description        |                                            |
| Task Execution ID  | <none>                                     |
| Root Execution ID  | <none>                                     |
| State              | RUNNING                                    |
| State info         | None                                       |
| Created at         | 2020-01-22 15:47:57                        |
| Updated at         | 2020-01-22 15:47:57                        |
+--------------------+--------------------------------------------+
```

In the earlier example output, the value of State is RUNNING. State should eventually read SUCCESS.

4. Re-check the value of State:

```
$ openstack workflow execution show edcf9103-e1a8-42f9-85c1-e505c055e0ed
```

```
+--------------------+--------------------------------------------+
| Field              | Value                                      |
+--------------------+--------------------------------------------+
| ID                 | edcf9103-e1a8-42f9-85c1-e505c055e0ed       |
| Workflow ID        | 8aa2ac9b-22ee-4e7d-8240-877237ef0d0a       |
| Workflow name      | tripleo.plan_management.v1.rotate_passwords |
| Workflow namespace |                                            |
| Description        |                                            |
| Task Execution ID  | <none>                                     |
| Root Execution ID  | <none>                                     |
| State              | SUCCESS                                    |
| State info         | None                                       |
| Created at         | 2020-01-22 15:47:57                        |
| Updated at         | 2020-01-22 15:48:39                        |
+--------------------+--------------------------------------------+
```

5. When the value of State is SUCCESS, you can retrieve passwords:

```
$ openstack workflow execution output show edcf9103-e1a8-42f9-85c1-
e505c055e0ed
```

You should see output similar to the following:

```
  {
     "status": "SUCCESS",
     "message": {
        "AdminPassword": "FSn0sS1aAHp8YK2fU5niM3rxu",
        "AdminToken": "dTP0Wdy7DtblG80M54r4a2yoC",
        "AodhPassword": "fB5NQdRe37BaBVEWDHVuj4etk",
        "BarbicanPassword": "rn7yk7KPafKw2PWN71MvXpnBt",
        "BarbicanSimpleCryptoKek": "lrC3sGlV7-D7-
V_PI4vbDfF1Ujm5OjnAVFcnihOpbCg=",
        "CeilometerMeteringSecret": "DQ69HdlJobhnGWoBC0jM3drPF",
        "CeilometerPassword": "qI6xOpofuiXZnG95iUe8Oxv5d",
        "CephAdminKey": "AQDGVPpdAAAAABAAZMP56/VY+zCVcDT81+TOjg==",
        "CephClientKey": "AQDGVPpdAAAAABAAanYtA0ggpcoCbS1nLeDN7w==",
        "CephClusterFSID": "141a5ede-21b4-11ea-8132-52540031f76b",
        "CephDashboardAdminPassword":
"AQDGVPpdAAAAABAAKhsx630YKDhQrocS4o4KzA==",
        "CephGrafanaAdminPassword":
"AQDGVPpdAAAAABAAKBojG+CO72B0TdBRR0paEg==",
        "CephManilaClientKey":
"AQDGVPpdAAAAABAAA1TVHrTVCC8xQ4skG4+d5A=="
     }
  }
(BZ#1600967)
```

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

- Nova-compute ironic driver tries to update BM node while the node is being cleaned up. The cleaning takes approximately five minutes but nova-compute attempts to update the node for approximately two minutes. After timeout, nova-compute stops and puts nova instance into ERROR state.
  As a workaround, set the following configuration option for nova-compute service:

```
[ironic]
api_max_retries = 180
```

As a result, nova-compute continues to attempt to update BM node longer and eventually succeeds. (BZ#1647005)

- There is a known issue where Red Hat OpenStack Platform 16.0, where metadata information required for configuring OpenStack instances is not available, and they might be started without connectivity.
  An ordering issue causes the haproxy wrappers not to be updated for the ovn_metadata_agent.

  A possible workaround is for the cloud operator to run the following Ansible command to restart the ovn_metadata_agent on select nodes after the update, to ensure that the ovn_metadata_agent is using an updated version of the haproxy wrapper script:

  **ansible -b <nodes> -i /usr/bin/tripleo-ansible-inventory -m shell -a "status=`sudo systemctl is-active tripleo_ovn_metadata_agent**; if test \"$status\" == \"active\"; then sudo systemctl restart tripleo_ovn_metadata_agent; echo restarted; fi"`

  In the earlier Ansible command, **nodes** may be a single node (for example, **compute-0**), all computes (for example, **compute***) or **"all"**.

  As the ovn_metadata_agent is most commonly found on compute nodes, the following Ansible command restarts the agent for all compute nodes in the cloud:

  **ansible -b compute* -i /usr/bin/tripleo-ansible-inventory -m shell -a "status=`sudo systemctl is-active tripleo_ovn_metadata_agent**; if test \"$status\" == \"active\"; then sudo systemctl restart tripleo_ovn_metadata_agent; echo restarted; fi"`

  After you restart the ovn_metadata_agent services, they use the updated haproxy wrapper script, which enables them to provide metadata to VMs when they are started. Affected VMs already running should behave normally when they are restarted after the workaround has been applied. (BZ#1790467)

- Red Hat OpenStack Platform 16.0 director, now supports multi-compute cell deployments. With this enhancement, your cloud is better positioned for scaling out, because each individual cell has its own database and message queue on a cell controller and reduces the load on the central control plane. For more information, see "Scaling deployments with Compute cells" in the "Instances and Images" guide. (BZ#1328124)

- Starting with this update, OSP deployments have full encryption between all the OVN services. All OVN clients (ovn-controller, neutron-server and ovn-metadata-agent) now connect to the OVSDB server using Mutual TLS encryption. (BZ#1601926)

- In Red Hat OpenStack Platform 16.0, a technology preview has been added to the Orchestration service (heat) for rsyslog changes:

  - Rsyslog is configured to collect and forward container logs to be functionally equivalent to the fluentd installation.

  - Administrators can configure rsyslog log forwarding in the same way as fluentd. (BZ#1623152)

- In Red Hat OpenStack Platform 16.0, you can now use director to specify an availability zone for the Block Storage service (cinder) back end type. (BZ#1700396)

- Red Hat OpenStack Platform director now enables you to deploy an additional node that can be used to add additional Bare Metal Provisioning conductor service resources for system provisioning during deployments. (BZ#1710093)

- In Red Hat OpenStack Platform 16.0, the Elastic Compute Cloud (EC2) API is no longer supported. The EC2 API support is now deprecated in director and will be removed in a future RHOSP release. (BZ#1754560)

- In Red Hat OpenStack Platform 16.0, the controller-v6.yaml file is removed. The routes that were defined in controller-v6.yaml are now defined in controller.yaml. (The controller.yaml file is a NIC configuration file that is rendered from values set in roles_data.yaml.)
  Previous versions of Red Hat OpenStack director, included two routes: one for IPv6 on the External network (default) and one for IPv4 on the Control Plane.

  To use both default routes, make sure that the controller definition in roles_data.yaml contains both networks in default_route_networks (for example, **default_route_networks: ['External', 'ControlPlane']**). (BZ#1631508)

- In Red Hat OpenStack Platform 16.0, you can now add custom Red Hat Ceph Storage configuration settings to any section of ceph.conf. Previously, custom settings were allowed only in the **[global]** section of ceph.conf. (BZ#1666973)

- In Red Hat OpenStack Platform 16.0, a new Orchestration service (heat) deployment parameter is available that enables administrators to turn on the nova metadata service on cell controllers:

  ```
  parameter_defaults:
    NovaLocalMetadataPerCell: True
  ```

  This new parameter automatically directs traffic from the OVN metadata agent on the cell computes to the nova metadata API service hosted on the cell controllers.

  Depending on the RHOSP topology, the ability to run the metadata service on cell controllers can reduce the traffic on the central control plane. (BZ#1689816)

- In Red Hat OpenStack Platform 16.0, a technology preview has been added to the Orchestration service (heat). A new parameter, **NovaSchedulerQueryImageType,** has been added that controls the Compute service (nova) placement and scheduler components query placement for image type (scheduler/query_placement_for_image_type_support).
  When set to true (the default), **NovaSchedulerQueryImageType** excludes compute nodes that do not support the disk format of the image used in a boot request.

  For example, the libvirt driver uses Red Hat Ceph Storage as an ephemeral back end, and does not support qcow2 images (without an expensive conversion step). In this case, enabling **NovaSchedulerQueryImageType** ensures that the scheduler does not send requests to boot a qcow2 image to compute nodes that use Red Hat Ceph Storage. (BZ#1710634)

Changes to the puppet-tripleo component:

- Red Hat OpenStack Platform director now offers a technology preview for fence_redfish, a fencing agent for the Redfish API. (BZ#1699449)

Changes to the python-django-horizon component:

- In the Red Hat OpenStack Platform 16.0 dashboard (horizon), there is now a new form for changing a user's password. This form automatically appears when a user tries to sign on with an expired password. (BZ#1628541)

Changes to the python-networking-ansible component:

- In Red Hat OpenStack Platform 16.0, a technology preview is added to the OpenStack Bare Metal service (ironic) to configure ML2 networking-ansible functionality with Arista Extensible Operating System (Arista EOS) switches. For more information, see "Enabling networking-ansible ML2 functionality," in the Bare Metal Provisioning guide. (BZ#1621701)

- In Red Hat OpenStack Platform 16.0, a technology preview has been added to modify switch ports to put them into trunking mode and assign more than one VLAN to them. (BZ#1622233)

Changes to the python-networking-ovn component:

- In Red Hat OpenStack Platform 16.0, live migrations with OVN enabled now succeed, as the flag, **live_migration_wait_for_vif_plug,** is enabled by default.
  Previously, live migrations failed, because the system was waiting for OpenStack Networking (neutron) to send **vif_plugged** notifications. (BZ#1716335)

- Currently, the OVN load balancer does not open new connections when fetching data from members. The load balancer modifies the destination address and destination port and sends request packets to the members. As a result, it is not possible to define an IPv6 member while using an IPv4 load balancer address and vice versa. There is currently no workaround for this issue. (BZ#1734301)

Changes to the python-novajoin component:

- Previously, when Novajoin lost its connection to the IPA server, it would immediately attempt to reconnect. Consequently, timing issues could arise and prevent the connection from being re-established. With this update, you can use **retry_delay** to set the number of seconds to wait before retrying the IPA server connection. As a result, this is expected to help mitigate the timing issues. (BZ#1767481)

Changes to the python-oslo-utils component:

- Previously, the regular expression for the oslo.util library was not updated, and it failed to recognize the output format from a newer version of the emulator, qemu (version 4.1.0). This fix in Red Hat OpenStack 16.0 updates the regular expression, and the oslo.util.imageutils library now functions properly. (BZ#1758302)

Changes to the python-tripleoclient component:

- Director has added the **overcloud undercloud minion install** command that you can use to configure an additional host to augment the Undercloud services. (BZ#1710089)

- Director now provides the ability to deploy an additional node that you can use to add additional heat-engine resources for deployment related actions. (BZ#1710092)

- In Red Hat OpenStack Platform 16.0, you are now able to push, list, delete, and show (show metadata) images on the local registry.

  - To push images from remote repository to the main repository:

    ```
    $ sudo openstack tripleo container image push docker.io/library/centos
    ```

  - To list the contents of the repository:

    ```
    $ openstack tripleo container image list
    ```

- To delete images:

```
$ sudo openstack tripleo container image delete
```

- To show metadata for an image:

```
$ openstack tripleo container image show (BZ#1545855)
```

- With this enhancement, overcloud node deletion requires user confirmation before the action will be performed to reduce the likelihood that the action is performed unintentionally. The **openstack overcloud node delete <node>** command requires a Y/n confirmation before the action executes. You can bypass this by adding **--yes** to the command line. (BZ#1593057)

4.2. RHBA-2020:2114 – RED HAT OPENSTACK PLATFORM 16.0.2 ADVISORY

The bugs contained in this section are addressed by advisory 2020:2114. Further information about this advisory is available at link: https://access.redhat.com/errata/RHBA-2020:2114.html.

Changes to the openstack-tripleo-common component:

- This update fixes authentication timeouts caused by slow transfer of container images. Previously, undercloud and overcloud pulls against container sources that require authentication could fail, and generate a 401 error, if the image transfer exceeded five minutes. Now, if the container fetching process exceeds 5 minutes, the code attempts to re-authenticate, preventing the timeout. (BZ#1813520)

- A Grafana Ceph 4.1 dependency causes Ceph dashboard bugs. The Ceph dashboard requires Ceph 4.1 and a Grafana container based on ceph4-rhel8. Presently, Red Hat supports ceph3-rhel7.3. This discrepancy causes the following dashboard bugs:

  - When you navigate to **Pools** > **Overall Performance**, Grafana returns the following error:

```
TypeError: l.c[t.type] is undefined
true
```

  - When you view a pool's performance details (**Pools** > select a pool from the list > **Performance Details**) the Grafana bar is displayed along with other graphs and values, but it should not be there.
    These bugs will be fixed after rebasing to a newer Grafana version. (BZ#1824093)

- This update fixes a bug that caused the **upload-puppet-modules** command to fail after the first invocation. A recent OpenStack command line interface update changed the OpenStack formats JSON data. That new format broke a script responsible for maintaining an internal URL used by the 'upload-puppet-modules' command. The script has been fixed to correctly handle the JSON data. Now the 'upload-puppet-modules' command functions correctly every time. (BZ#1808369)

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

- There is a known issue for the Red Hat OpenStack Platform Load-balancing service: the containers octavia_api and octavia_driver_agent fail to start when rebooting a node.
  The cause for this issue is that the directory, /var/run/octavia, does not exist when the node is rebooted.

To fix this issue, add the following line to the file, /etc/tmpfiles.d/var-run-octavia.conf:

```
d /var/run/octavia 0755 root root - -
```

(BZ#1795956)

- RHOSP 16.0 works only with RHEL 8.1. Ensure that all the hosts of your OSP deployment are pinned to RHEL 8.1 before running the update.
  See "Locking the environment to a Red Hat Enterprise Linux release" [1] in the guide "Keeping Red Hat OpenStack Platform Updated."

  [1] https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.0/html/keeping_red_hat_openstack_platform_updated/prepar for-a-minor-update#locking-the-environment-to-a-red-hat-enterprise-linux-release
  (BZ#1823835)

- This update fixes a bug that prevented display of Granfana layouts in Ceph dashboard iframes in high availability scenarios. Previously, the Grafana fronend could not be reached on the storage network. GET requests got stuck. This fix moves the Grafana frontend to the same network used by the Ceph dashboard. Now the GET requests succeed and the Grafana layouts are available in the Ceph dashboard. (BZ#1815037)

- Director operations involving the RADOS gateway no longer require interaction with puppet-ceph. Previously, tripleo-heat-templates had a dependency on puppet-ceph for the RADOS gateway component deployment. The move to tripleo-ansible eliminates this dependency. (BZ#1695898)

Changes to the python-tripleoclient component:

- This update fixes a bug that caused the **openstack overcloud node import** command to fail with iPXE disabled ('ipxe_enabled=False'). With iPXE disabled, you must use the **--http-boot** argument to specify the location of kernel and ramdisk images (**--http-boot /var/lib/ironic/tftpboot**). Previously, the **openstack overcloud node import** command ignored the **--http-boot** argument. The nodes failed to deploy. Now the **openstack overcloud node import** command responds correctly to the **http-boot** argument and the nodes are deployed as expected. (BZ#1793175)