# Red Hat OpenStack Platform 15

# Release Notes

Release details for Red Hat OpenStack Platform 15

# Red Hat OpenStack Platform 15 Release Notes

Release details for Red Hat OpenStack Platform 15

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

## 1.1. ABOUT THIS RELEASE

This release of Red Hat OpenStack Platform is based on the OpenStack "Stein" 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 "Stein" release itself are available at the following location: https://releases.openstack.org/stein/index.html.

Red Hat OpenStack Platform uses components from other Red Hat products. See the following link(s) for specific information pertaining to the support of these components: 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. See the following link for more details about the add-on: http://www.redhat.com/products/enterprise-linux-add-ons/high-availability/. See the following link for details about the package versions to use in combination with Red Hat OpenStack Platform: https://access.redhat.com/site/solutions/509783.

## 1.2. REQUIREMENTS

This version of Red Hat OpenStack Platform runs on Red Hat Enterprise Linux version 8.2.

The Red Hat OpenStack Platform dashboard is a web-based interface that allows you to manage OpenStack resources and services. The dashboard for this release supports the latest stable versions of the following web browsers:

- Chrome

- Mozilla Firefox

- Mozilla Firefox ESR

- Internet Explorer 11 and later (with **Compatibility Mode** disabled)

> **NOTE**
>
> Prior to deploying Red Hat OpenStack Platform, it is important to consider the characteristics of the available deployment methods. For more information, refer to the 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 15.

You can install Red Hat OpenStack Platform 15 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

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) | **rhel-8-for-x86_64-baseos-rpms** | Base operating system repository for x86_64 systems. |
| Red Hat Enterprise Linux 8 for x86_64 – AppStream (RPMs) | **rhel-8-for-x86_64-appstream-rpms** | Contains Red Hat OpenStack Platform dependencies. |
| Red Hat Enterprise Linux 8 for x86_64 – High Availability (RPMs) | **rhel-8-for-x86_64-highavailability-rpms** | High availability tools for Red Hat Enterprise Linux. 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 15 for RHEL 8 (RPMs) | **openstack-15-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 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 15 for RHEL 8 (RPMs) | **openstack-15-for-rhel-8-ppc64le-rpms** | Core Red Hat OpenStack Platform repository for ppc64le systems. |

## 1.10.2. Overcloud repositories

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) | **rhel-8-for-x86_64-baseos-rpms** | Base operating system repository for x86_64 systems. |
| Red Hat Enterprise Linux 8 for x86_64 – AppStream (RPMs) | **rhel-8-for-x86_64-appstream-rpms** | Contains Red Hat OpenStack Platform dependencies. |
| Red Hat Enterprise Linux 8 for x86_64 – High Availability (RPMs) | **rhel-8-for-x86_64-highavailability-rpms** | High availability tools for Red Hat Enterprise Linux. 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. |

| Name | Repository | Description of Requirement |
| --- | --- | --- |
| 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 15 for RHEL 8 (RPMs) | **openstack-15-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. This repository should be enabled for all Compute nodes targeted for RT-KVM. NOTE: You need a separate subscription to a **Red Hat OpenStack Platform for Real Time** SKU before you can 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. This repository should be enabled 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 before you can 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 15 for RHEL 8 (RPMs) | **openstack-15-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. RED HAT ENTERPRISE LINUX 8 FEATURES THAT AFFECT RED HAT OPENSTACK PLATFORM 15

This section outlines new features in Red Hat Enterprise Linux 8 that affect Red Hat OpenStack Platform 15.

Red Hat OpenStack Platform 15 now uses Red Hat Enterprise Linux 8 for the operating system. This includes the undercloud node, the overcloud nodes, and containerized services. Some key differences between Red Hat Enterprise Linux 7 and 8 impact the architecture of Red Hat OpenStack Platform 15. The following list provides information on these key differences and their effect on Red Hat OpenStack Platform:

**New Red Hat Enterprise Linux 8 repositories**

In addition to the Red Hat OpenStack Platform 15 repository, OpenStack Platform now uses a new set of repositories specific to Red Hat Enterprise Linux 8. This includes the following repositories:

- **BaseOS** for the main operating system packages.

- **AppStream** for dependencies such as Python 3 packages and virtualization tools.

- **High Availability** for Red Hat Enterprise Linux 8 versions of high availability tools.

- **Red Hat Ansible Engine** for the latest supported version of Ansible engine.

This change affects the repositories you must enable for both the undercloud and overcloud.

**Red Hat Enterprise Linux 8 container images**

All OpenStack Platform 15 container images use Red Hat Enterprise Linux 8 Universal Base Image (UBI) as a basis. OpenStack Platform director automatically configures these container images during undercloud and overcloud creation.

> **IMPORTANT**
>
> Red Hat does not support running OpenStack Platform containers based on Red Hat Enterprise Linux 7 on a Red Hat Enterprise Linux 8 host.

**Red Hat Enterprise Linux 8 bare metal images**

All OpenStack Platform 15 overcloud kernel, ramdisk, and QCOW2 images use Red Hat Enterprise Linux 8 as a basis. This includes OpenStack Bare Metal (ironic) introspection images.

**Python 3 packages**

All OpenStack Platform 15 services use Python 3 packages.

**New Container Tools**

Red Hat Enterprise Linux 8 no longer includes Docker. Instead, Red Hat provides new tools for building and managing containers:

- **Pod Manager (Podman)** is a container management tool that implements almost all Docker CLI commands, not including commands related to Docker Swarm. Podman manages pods,

containers, and container images. One of the major differences between Podman and Docker is Podman can manage resources without a daemon running in the background. For more information on Podman, see the Podman website.

- **Buildah** specializes in building Open Containers Initiative (OCI) images, which you use in conjunction with Podman. Buildah commands replicate what you find in a Dockerfile. Buildah also provides a lower-level **coreutils** interface to build container images, which allows you to build containers without requiring a Dockerfile. Buildah can also use other scripting languages to build container images without requiring a daemon.

### Docker Registry Replacement

Red Hat Enterprise Linux 8 no longer includes the **docker-distribution** package, which installed a Docker Registry v2. To maintain the compatibility, OpenStack Platform 15 includes an Apache web server and a virtual host called **image-serve**, which provides a container registry. Like **docker-distribution**, this registry uses port 8787/TCP with SSL/TLS disabled.
This registry is a **read-only container registry** and does not support **podman push** nor **buildah push** commands. This means the registry does not allow you to push non-director and non-OpenStack Platform containers. However, you can modify supported OpenStack Platform images with the director's container preparation workflow, which uses the **ContainerImagePrepare** parameter.

### Network Time Synchronization

Red Hat Enterprise Linux 8 no longer includes **ntpd** to synchronize your system clock. Red Hat Enterprise Linux 8 now provides **chronyd** as a replacement service. The director configures **chronyd** automatically but be advised that manual time synchronization requires the execution of the **chronyc** client.

### High Availability and Shared Services

- **Pacemaker 2.0 support.** This release upgrades the version of Pacemaker to 2.0 to support deployment on top of Red Hat Enterprise Linux 8, including support for Knet and multiple NICs. You can now use the director to configure fencing with Pacemaker for your high availability cluster.

- **HAProxy 1.8 support in director.** This release upgrades the version of HAProxy to 1.8 to support deployment on top of Red Hat Enterprise Linux 8. You can now use the director to configure HAProxy for your high availability cluster.

## 2.2. RED HAT OPENSTACK PLATFORM DIRECTOR

This section outlines the top new features for the director.

### Deploy an all-in-one overcloud

This release includes the capability to deploy standalone overclouds that contain Controller and Compute services. (High availability is not supported.)

- Use the **Standalone.yaml** role file to deploy all-in-one overclouds.

- Enable and disable services on the all-in-one overcloud

- Customize services with additional environment files.

- The following services are enabled by default:

  - Keystone

- Nova and related services

- Neutron and related services

- Glance

- Cinder

- Swift

- Horizon

Unified composable service templates

This release includes a unified set of composable service templates in **/usr/share/openstack-tripleo-heat-templates/deployment/**. These templates merge service configuration from previous template sets:

- The containerized templates previously in **/usr/share/openstack-tripleo-heat-templates/docker/services/**.

- The Puppet-based templates previously in **/usr/share/openstack-tripleo-heat-templates/puppet/services/**.

Most service resources, which start with the **OS::TripleO::Services** Heat namespace, now refer to the unified template set. A few minor services still refer to legacy templates in **/usr/share/openstack-tripleo-heat-templates/puppet/services/**.

## 2.3. COMPUTE

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

NUMA-aware vSwitches (full support)

OpenStack Compute now takes into account the NUMA node location of physical NICs when launching a Compute instance. This helps to reduce latency and improve performance when managing DPDK-enabled interfaces.

Parameters to control VM status after restarting Compute nodes

The following parameters are now available to determine whether to resume the previous state of virtual machines (VMs) after you restart Compute nodes:

- **NovaResumeGuestsStateOnHostBoot** (True/False)

- **NovaResumeGuestsShutdownTimeout** (default 300s)

Using placement service traits to schedule hosts

Compute nodes now report the CPU instruction set extensions supported by the host, as traits against the resource provider corresponding to that Compute node. An instance can specify which of these traits it requires using image metadata or flavor extra specs. This feature is available when your deployment uses a Libvirt QEMU (x86), Libvirt KVM (x86), or Libvirt KVM (ppc64) virtualization driver. For more information, see Configure Placement Service Traits.

## 2.4. OPENSTACK NETWORKING

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

ML2/OVN Default Plugin

OVN replaces OVS as the default ML2 mechanism driver. OVN offers a shared networking back end across the RedHat portfolio, giving operators a consistent experience across multiple products. Its architecture offers simpler foundations than OVS, and better performance.

## 2.5. STORAGE

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

### Deprecation of Data Processing service (sahara)

In Red Hat OpenStack Platform 15, the Data Processing service (sahara) is deprecated. In a future RHOSP release, the the Data Processing service will be removed.

### Simultaneously attach a volume to multiple instances

In Red Hat OpenStack Platform 15, if the back end driver supports it, you can now simultaneously attach a volume to multiple instances for both the Block Storage service (cinder) and the Compute service (nova). This feature addresses the use case for clustered application workloads that typically requires active/active or active/standby scenarios.

### Image Service - Deploy glance-cache on far-edge nodes

This feature enables deploying Image service (glance) caching at the edge. This feature improves the boot time for instances and decreases the bandwidth usage between core and edge sites, which is useful in medium to large edge sites to avoid using compute nodes to fetch images from the core site.

## 2.6. 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 15 Beta, director deploys Red Hat Ceph Storage 4 Beta. Red Hat OpenStack Platform 15 running on RHEL8 enables you to connect to a preexisting external Red Hat Ceph Storage 3 cluster running on RHEL7.

## 2.7. TECHNOLOGY PREVIEWS

This section outlines features that are in technology preview in Red Hat OpenStack Platform 15.

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

### 2.7.1. 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.

### 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 on Ceph RADOS Block Device (RBD) back ends only, if the back end driver supports this 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 15, 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.

### Using separate heat stacks

You can now use separate heat stacks for different types of nodes. For example, you can have a stack just for the control plane, a stack for Compute nodes, and another stack for Hyper Converged Infrastructure (HCI) nodes. This approach has the following benefits:

- Management – You can modify and manage your nodes without needing to make changes to the control plane stack.

- Scaling out – You do not need to update all nodes just to add more Compute or Storage nodes; the separate heat stack means that these operations are isolated to selected node types.

- Edge sites – You can segment an edge site within its own heat stack, thereby reducing network and management dependencies on the central data center. The edge site must have its own Availability Zone for its Compute and Storage nodes.

### 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 to store the memory backing file in 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.

**Support added for templated cell mapping URLs**

Director now provides cell mapping URLs for the database and message queue URLs as templates, by using variables to represent values such as the username and password. The following properties, defined in the Compute configuration file of the director, specify the variable values:

- **database_connection**: **[database]/connection**

- **transport_url**: **[DEFAULT]/transport_url**

**Support added for configuring the maximum number of disk devices that can be attached to a single instance**

A new Compute (nova) parameter is available, **max_disk_devices_to_attach**, which specifies the maximum number of disk devices that can be attached to a single instance. The default is unlimited (-1). The following example illustrates how to change the value of **max_disk_devices_to_attach** to "30":

```
parameter_defaults:
  ComputeExtraConfig:
    nova::config::nova_config:
      [compute]/max_disk_devices_to_attach:
        value: '"30"'
```

# 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 15 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. Enhancements

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

BZ#1240852

> In Red Hat OpenStack Platform 15, you can specify MTU (maximum transmission unit) settings for each network, and RHOSP will automatically write those settings to the network interface configuration templates. MTU values should be set in the network_data.yaml file.
>
> This enhancement alleviates the step of manually updating the network templates for each role, and reduces the likelihood of manual entry errors.

BZ#1484601

> The Shared File Systems service (manila) API now supports Transport Layer Security (TLS) endpoints on the internal API network, through SSL/TLS certificates. The Shared File Systems service is automatically secured when you opt to secure Red Hat OpenStack Platform during deployment.

BZ#1535066

> In Red Hat OpenStack Platform 15, which depends on Red Hat Enterprise Linux 8, there is a new default Time service, chrony.
>
> With this switch, Red Hat highly recommends that you use multiple Network Time Protocol (NTP) servers for both the undercloud and overcloud deployments.

BZ#1547728

> In Red Hat Open Stack Platform 15, the Data Processing service (sahara) plug-ins have been decoupled and are now installed as libraries.
>
> To obtain newer versions of Data Processing service plug-ins, you no longer have to upgrade RHOSP. Instead, install the newest version of the desired plug-in.

BZ#1585012

> You can now configure automatic restart of VM instances on a Compute node if the compute node reboots without first migrating the instances.

With the following two new parameters, you can configure the Red Hat OpenStack Platform Compute service (nova) and the libvirt-guests agent to shut down VM instances gracefully and start them when the Compute node reboots:

- NovaResumeGuestsStateOnHostBoot (True or False)
- NovaResumeGuestsShutdownTimeout (default, 300s)

BZ#1619762

In Red Hat OpenStack Platform 15, director uses version 5.5 of Puppet.

BZ#1626139

In Red Hat OpenStack Platform 15, a new role and environment file have been added to enable the undercloud to deploy an all-in-one overcloud node that contains both the controller services and compute services. The new role and the new environment file are named, respectively, roles/Standalone.yaml and environments/standalone/standalone-overcloud.yaml.

Because this new architecture does not yet support high availability, Red Hat cannot guarantee zero down time during RHOSP 15 updates and upgrades. For this reason, Red Hat highly recommends that you properly back up your system.

BZ#1633146

Red Hat OpenStack Platform director now has the ability to control Block Storage service (cinder) snapshots on NFS back ends. A new director parameter, CinderNfsSnapshotSupport, has a default value of True.

BZ#1635862

Using the Red Hat OpenStack Platform director, you can now configure the Image service (glance) to have an optional local image cache. You turn on the image cache, by setting the "GlanceCacheEnabled" property to True.

A typical use case for the image cache is edge computing. Because the Image service resides at central site, you can deploy and enable the image cache at remote sites and save bandwidth and reduce the Image service's boot time.

BZ#1647057

With Paunch you can now manage container memory consumption using three new attributes: mem_limit, memswap_limit, and mem_swappiness.

BZ#1661022

In Red Hat OpenStack Platform 15, if the back end driver supports it, you can now simultaneously attach a volume to multiple machines for both the Block Storage service (cinder) and the Compute service (nova). This feature addresses the use case for clustered application workloads that typically requires active/active or active/standby scenarios.

BZ#1666529

In Red Hat OpenStack Platform 15, the Image service (glance) is automatically configured for any glance-import execution to convert imported images into RAW format when Red Hat Ceph Storage is used as the back end for the Image service.

BZ#1693268

The Load Balancing service (octavia) now provides the capability to refine access policies for its load balancers, by allowing you to change security group ownership to a security group associated with a user project. (The user project must be on the whitelist.)

In previous RHOSP releases, you could not restrict access to the load balancer, because octavia exclusively assigned the project ID to the security group associated with the VIP and VRRP ports on the load balancing agent (amphora).

## 3.1.2. 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#1466008

The director can now deploy different, isolated Ceph clusters into different Edge zones by creating an overcloud composed of multiple Heat stacks. For example, the director can deploy an overcloud consisting of a Heat stack for the control plane (Controller nodes) and multiple additional stacks for Edge zones (Computes and Ceph Storage nodes or Compute and HCI nodes).

BZ#1504662

Neutron bulk port creation (create multiple ports in a single request) has been optimized for speed and is now significantly faster. The benefits of this improvement include faster initialization of containers via Kuryr on neutron networks.

BZ#1526109

A new Red Hat OpenStack Platform Bare Metal service (ironic) driver for XClarity managed Lenovo devices is available. The xclarity driver provides more reliable operation on Lenovo devices managed with XClarity, and opportunities for additional vendor-specific features in the future.

BZ#1593758

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 15, 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.

BZ#1601576

Red Hat OpenStack Platform undercloud networks are now layer 3 (L3) capable. This enhancement enables all segments to use one network, and alleviates the need for service net map overrides.

This enhancement is important for Red Hat OpenStack Platform edge computing sites that deploy roles in different sites and make service net map overrides unwieldy.

BZ#1624486

As a technology preview in Red Hat OpenStack Platform 15, the novajoin service tech uses the new, versioned format of notifications sent by the Compute service (nova).

To enable the new format, set the value of the new configuration setting, configuration_format, to "versioned." The default value for configuration_format is "unversioned".

In a future version of RHOSP, unversioned notifications will be deprecated.

BZ#1624488

As a technology preview in Red Hat OpenStack Platform 15, the novajoin service uses the Python 3 runtime.

BZ#1624490

With this technology preview, it is possible to configure Barbican through Director to store secrets using the ATOS Trustway Proteccio NetHSM. This is mediated through the Barbican PKCS#11 back-end plugin.

The technology preview is provided in the following packages:
 - openstack-barbican
 - tripleo-heat-templates

BZ#1624491

With this technology preview, it is possible to configure Barbican through director to store secrets using the nCipher NetShield Connect NetHSM. This is mediated through the Barbican PKCS#11 back end plug-in.

The technology preview is provided in the following packages:
 - openstack-barbican
 - tripleo-heat-templates

BZ#1636040

With Red Hat OpenStack Platform director you can now deploy the Block Storage service (cinder) in an active-active configuration on Ceph RADOS Block Device (RBD) back ends only.

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).

BZ#1636179

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.

BZ#1740715

Because Red Hat Ceph Storage 4 is at beta when Red Hat OpenStack Platform 15 is at GA, a new configuration option has been added to RHOSP 15 to prevent any accidental deployments of Red Hat Ceph Storage 4 Beta in a production environment.

The new Orchestration service (heat) configuration option, EnableRhcs4Beta, is set by default to "False", and therefore prevents director from deploying Red Hat Ceph Storage 4 Beta by accident.

### 3.1.3. 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#1585835

The Shared File Systems service (manila) API now runs behind the Apache HTTP Server (httpd). The Apache error and access logs from the Shared File Systems service are available in /var/log/containers/httpd/manila-api on all the nodes that run the manila API container.

The log location of the main API service (manila-api) has not changed, and continues to be written on each node in /var/log/containers/manila/.

BZ#1613038

The Block Storage service (cinder) command, "snapshot-manageable-list," now lists the snapshots on the back end for Red Hat Ceph RADOS block devices (RBD).

BZ#1689913

In Red Hat OpenStack Platform 15, the director parameter used during overcloud container preparation, deltarpm, has been renamed to, drpm.

BZ#1722036

Because Red Hat Ceph Storage 4 is at beta when Red Hat OpenStack Platform 15 is at GA, a new configuration option has been added to RHOSP 15 to prevent any accidental deployments of Red Hat Ceph Storage 4 Beta in a production environment.

The new Orchestration service (heat) configuration option, EnableRhcs4Beta, is set by default to "False", and therefore prevents director from deploying Red Hat Ceph Storage 4 Beta by accident.

BZ#1730689

There is a known issue wherein deployments will fail with the following message.

`puppet-user: Error: Parameter value failed on Vs_config[other_config:n-revalidator-threads]: Invalid external_ids 1. Requires a String, not a Integer`

This is due to tripleo parameters of type integer being expected by puppet to be of type string. To work around, include the following in deployment templates:

ComputeOvsDpdkSriovExtraConfig:
  "vswitch::dpdk::handler_cores": "1"
  "vswitch::dpdk::revalidator_cores": "1"

BZ#1743701

In Red Hat OpenStack Platform 15, director can only deploy Red Hat Ceph Storage v4. At this time, Ceph Storage v4 is still in its beta version. OpenStack Platform 15 will not support director-deployed Ceph until Ceph Storage v4 is generally available.

For testing purposes, you can deploy Ceph Storage v4 beta, but the beta version is not supported for use in production. Refer to the documentation for instructions on how to enable Ceph Storage v4 beta.

## 3.1.4. Known Issues

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

BZ#1543414

When running Red Hat OpenStack Platform 15 on a Q35 machine, there is a maximum limit of 500 devices. This is known problem with QEMU, an open source virtualizer and machine emulator.

BZ#1697335

When running the command "openstack stack show <stack_name>" on a stack with a large amount of data (for example, the 'overcloud' stack), the output can be difficult to read because some columns are too wide.

Red Hat recommends that you change the default output width.

Here is an example:

$ openstack stack show overcloud --max-width 100

BZ#1713329

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 the Red Hat OpenStack Platform director.

BZ#1741244

Red Hat OpenStack Platform (RHOSP) does not yet support upgrading to version 15 from earlier RHOSP versions. Support for upgrading will be added to a future update of RHOSP 15.

BZ#1749443

The Compute services (nova) can fail to deploy because the nova_wait_for_compute_service script is unable to query the Nova API. If a remote container image registry is used outside of the undercloud, the Nova API service might not finish deploying in time.

The workaround is to rerun the deployment command, or to use a local container image registry on the undercloud.

BZ#1751942

If you use Security Group rules that span across a port range (--dst-port X:Y), an OVN bug causes traffic filtering to fail and all traffic to be dropped.

Workaround: Create one rule per port instead of using a port range.

BZ#1752950

Currently, you cannot use Orchestration (heat) templates with the director to deploy an overcloud that requires NFS as an Image service (glance) back end. There is currently no workaround for this issue.

### 3.1.5. Deprecated Functionality

The items in this section are either no longer supported, or will no longer be supported in a future release.

BZ#1584213

In Red Hat OpenStack Platform 15, a part of the Telemetry service, gnocchi, has been deprecated.

In a future RHOSP version, gnocchi, and the rest of the Telemetry service will be removed and replaced by the Red Hat Service Assurance Framework.

BZ#1640962

In Red Hat OpenStack Platform 15, the Alarm service (aodh) that is part of the Telemetry service, is deprecated.

In a future Red Hat OpenStack Platform version, the Alarm service will be removed.

BZ#1663449

The OpenStack EC2 API is deprecated in this release and is no longer supported.

BZ#1676951

In Red Hat OpenStack Platform 15, the monitoring agent, Sensu client service, is deprecated.

In a future Red Hat OpenStack Platform version, the Sensu client service will be removed.

BZ#1686583

> In Red Hat OpenStack Platform 15, the Data Processing service (sahara) is deprecated, and will be removed in version 16. Support for the Data Processing service continues in Red Hat OpenStack Platform 15 and earlier supported versions.

BZ#1702694

> In Red Hat OpenStack Platform 15, Red Hat OpenStack director (TripleO) no longer supports deploying Red Hat OpenShift Container Platform 3.11 clusters on bare metal nodes using the OpenShift installation playbooks (provided in the openshift-ansible package) and Orchestration service (heat) templates.
>
> To deploy OpenShift 3.11 on bare metal nodes, use the OpenShift installation playbooks exclusively without Orchestration service templates. You can provision Red Hat Enterprise Linux on bare metal nodes using Red Hat OpenStack Platform with the Bare Metal service (ironic) or by performing a manual installation.

BZ#1722809

> In Red Hat OpenStack Platform 15, the legacy network scripts are deprecated. In a future Red Hat OpenStack Platform version, the legacy network scripts will be removed and replaced by Red Hat Enterprise Linux NetworkManager.

BZ#1752660

> In Red Hat OpenStack Platform 15, the Nova vCenter plug-in is deprecated. It will be removed in version 16.

## 3.2. RED HAT OPENSTACK PLATFORM 15 MAINTENANCE RELEASE - OCTOBER 3, 2019

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#1693268

> The Load Balancing service (octavia) now provides the capability to refine access policies for its load balancers, by allowing you to change security group ownership to a security group associated with a user project. (The user project must be on the whitelist.)
>
> In previous RHOSP releases, you could not restrict access to the load balancer, because octavia exclusively assigned the project ID to the security group associated with the VIP and VRRP ports on the load balancing agent (amphora).

### 3.2.2. 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#1504662

> Neutron bulk port creation (create multiple ports in a single request) has been optimized for speed and is now significantly faster. The benefits of this improvement include faster initialization of containers via Kuryr on neutron networks.

### 3.2.3. Known Issues

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

BZ#1776406

> The internal locking model in the container infrastructure changed between podman 1.0.5 (RHEL 8.0) and 1.4.2 (RHEL 8.1). As a result, you cannot update a Red Hat OpenStack Platform (RHOSP) 15 deployment from Red Hat Enterprise Linux (RHEL) 8.0 to RHEL 8.1 without downtime.
>
> There is currently no known workaround. Red Hat is actively working on a solution to this problem. New deployments on RHEL 8.1 with the RHEL 8.1 RHOSP 15 containers work correctly.
>
> If you deployed RHOSP 15 on RHEL 8.0, contact Red Hat support for advisement.

BZ#1776851

> Due to an incompatibility between the version of pacemaker in Red Hat Enterprise Linux (RHEL) 8.0-based containers and RHEL 8.1-based hosts, you cannot migrate from RHEL 8.0 to 8.1 for deployed Red Hat OpenStack Platform (RHOSP) 15 environments.
>
> There is currently no known workaround. Red Hat is actively working on a solution to this problem. New deployments of RHOSP 15 on RHEL 8.1 work correctly.
>
> If you deployed RHOSP 15 on RHEL 8.0, contact Red Hat Support for advisement.

## 3.3. RED HAT OPENSTACK PLATFORM 15 MAINTENANCE RELEASE – MARCH 4, 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#1618894

> In this update, Red Hat OpenStack director deploys Red Hat Ceph Storage 4.

BZ#1696658

> With this update, you can configure NUMA affinity for most neutron networks. This is useful to ensure that instances are placed on the same host NUMA node as the NIC that provides external

connectivity to the vSwitch. This feature is available for networks that use a 'provider:network_type' of 'flat' or 'vlan' and a 'provider:physical_network' (L2 networks) or networks that use a 'provider:network_type' of 'vxlan', 'gre' or 'geneve' (L3 networks).

**BZ#1751809**

With this update, the credentials that you supply in the **ContainerImageRegistryCredentials** parameter pass to ceph-ansible automatically if the registry name matches the registry name in the **ceph_namespace** parameter.

**BZ#1783354**

With this update, you can 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 configurable only on a host-level basis for PCI passthrough devices.

### 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#1466008**

Director can now deploy different, isolated Ceph clusters into different Edge zones by creating an overcloud composed of multiple heat stacks. For example, director can deploy an overcloud that consists of a heat stack for the control plane (Controller nodes) and multiple additional stacks for Edge zones (Computes and Ceph Storage nodes or Compute and HCI nodes).

### 3.3.3. Known Issues

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

**BZ#1749443**

The Compute services (nova) can fail to deploy because the nova_wait_for_compute_service script is unable to query the Nova API. If you use a remote container image registry outside of the undercloud, the Nova API service might not finish deploying in time.

The workaround is to rerun the deployment command, or to use a local container image registry on the undercloud.

**BZ#1764025**

There is a known issue that OVN load balancer does not open new connections while fetching data from members. Instead, the load balancer modifies destination address and destination port and sends request packets to the member. As a result, you cannot define IPv6 members if you use IPv4 load balancer addresses, and you cannot define IPv4 members if you use IPv5 load balancer addresses.

There is currently no workaround for this issue.

# CHAPTER 4. TECHNICAL NOTES

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

## 4.1. RHEA-2019:2811 — REDHAT OPENSTACK PLATFORM 15 GENERAL AVAILABILITY ADVISORY

The enhancements and bug fixes contained in this section are addressed by advisory RHEA-2019:2811. Further information about this advisory is available at link: https://access.redhat.com/errata/RHEA-2019:2811

Changes to the ansible-role-tripleo-modify-image component:

- In Red Hat OpenStack Platform 15, the director parameter used during overcloud container preparation, deltarpm, has been renamed to, drpm. (BZ#1689913)

Changes to the distribution component:

- Skydive is a network analysis service that was designated as Technology Preview support in Red Hat OpenStack Platform 14. In RHOSP 15, Skydive has been removed. (BZ#1749427)

Changes to the Networking component:

- In Red Hat OpenStack Platform 15, the Kuryr-Kubernetes container network interface (CNI) plug-in is highly available (active/passive mode). (BZ#1579371)

Changes to the openstack-barbican component:

- With this technology preview, it is possible to configure Barbican through Director to store secrets using the ATOS Trustway Proteccio NetHSM. This is mediated through the Barbican PKCS#11 back-end plugin.
  The technology preview is provided in the following packages: openstack-barbican tripleo-heat-templates (BZ#1624490)

- With this technology preview, it is possible to configure Barbican through Director to store secrets using the nCipher NetShield Connect NetHSM. This is mediated through the Barbican PKCS#11 back-end plugin.
  The Technology Preview is provided in the following packages: openstack-barbican tripleo-heat-templates (BZ#1624491)

Changes to the openstack-cinder component:

- In Red Hat OpenStack Platform 15, if the back end driver supports it, you can now simultaneously attach a volume to multiple machines for both the Block Storage service (cinder) and the Compute service (nova). This feature addresses the use case for clustered application workloads that typically requires active/active or active/standby scenarios. (BZ#1661022)

- The Block Storage service (cinder) command, "snapshot-manageable-list," now lists the snapshots on the back end for Red Hat Ceph RADOS block devices (RBD). (BZ#1613038)

Changes to the openstack-ironic component:

- A new Red Hat OpenStack Platform Bare Metal service (ironic) driver for XClarity managed Lenovo devices is available. The xclarity driver provides more reliable operation on Lenovo

devices managed with XClarity, and opportunities for additional vendor-specific features in the future. (BZ#1526109)

- 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 15, 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. (BZ#1593758)

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 the Red Hat OpenStack Platform director. (BZ#1713329)

Changes to the openstack-nova component:

- In earlier Red Hat OpenStack Platform versions, the RHOSP Compute service (nova) diagnostics command returned an "IndexError" on compute instances that used VFIO interfaces.
  In RHOSP 15, this problem has been addressed. The diagnostics command now retrieves interface data directly from the guest XML, and appropriately adds NICs to the diagnostics object. (BZ#1649688)

Changes to the openstack-sahara component:

- In Red Hat Open Stack Platform 15, the Data Processing service (sahara) plug-ins have been decoupled and are now installed as libraries.
  To obtain newer versions of Data Processing service plug-ins, you no longer have to upgrade RHOSP. Instead, install the newest version of the desired plug-in. (BZ#1547728)

Changes to the openstack-tripleo-common component:

- In Red Hat OpenStack Platform 15, Red Hat OpenStack director (TripleO) no longer supports deploying Red Hat OpenShift Container Platform 3.11 clusters on bare metal nodes using the OpenShift installation playbooks (provided in the openshift-ansible package) and Orchestration service (heat) templates.
  To deploy OpenShift 3.11 on bare metal nodes, use the OpenShift installation playbooks exclusively without Orchestration service templates. You can provision Red Hat Enterprise Linux on bare metal nodes using Red Hat OpenStack Platform with the Bare Metal service (ironic) or by performing a manual installation. (BZ#1702694)

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

- In Red Hat OpenStack Platform 15, which depends on Red Hat Enterprise Linux 8, there is a new default Time service, chrony.
  With this switch, Red Hat highly recommends that you use multiple Network Time Protocol (NTP) servers for both the undercloud and overcloud deployments. (BZ#1535066)

- You can now configure automatic restart of VM instances on a Compute node if the compute node reboots without first migrating the instances.

With the following two new parameters, you can configure the Red Hat OpenStack Platform Compute service (nova) and the libvirt-guests agent to shut down VM instances gracefully and start them when the Compute node reboots:

- NovaResumeGuestsStateOnHostBoot (True or False)

- NovaResumeGuestsShutdownTimeout (default, 300s) (BZ#1585012)

- The Shared File Systems service (manila) API now runs behind the Apache HTTP Server (httpd). The Apache error and access logs from the Shared File Systems service are available in /var/log/containers/httpd/manila-api on all the nodes that run the manila API container. The log location of the main API service (manila-api) has not changed, and continues to be written on each node in /var/log/containers/manila/. (BZ#1585835)

- Red Hat OpenStack Platform undercloud networks are now layer 3 (L3) capable. This enhancement enables all segments to use one network, and alleviates the need for service net map overrides.
  This enhancement is important for Red Hat OpenStack Platform edge computing sites that deploy roles in different sites and make service net map overrides unwieldy. (BZ#1601576)

- In Red Hat OpenStack Platform 15, a new role and environment file have been added to enable the undercloud to deploy an all-in-one overcloud node that contains both the controller services and compute services. The new role and the new environment file are named, respectively, roles/Standalone.yaml and environments/standalone/standalone-overcloud.yaml. Because this new architecture does not yet support high availability, Red Hat cannot guarantee zero down time during RHOSP 15 updates and upgrades. For this reason, Red Hat highly recommends that you properly back up your system. (BZ#1626139)

- Using the Red Hat OpenStack Platform director, you can now configure the Image service (glance) to have an optional local image cache. You turn on the image cache, by setting the "GlanceCacheEnabled" property to True.
  A typical use case for the image cache is edge computing. Because the Image service resides at central site, you can deploy and enable the image cache at remote sites and save bandwidth and reduce the Image service's boot time. (BZ#1635862)

- 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. (BZ#1636179)

- Previously, when using TLS Everywhere, your controller node was required to access IdM through the **ctlplane** network. As a result, if traffic was routed through a different network, then the overcloud deployment process would fail due to **getcert** errors. To address this, IdM enrolment has been moved into a composable service that runs within **host_prep_tasks**; this runs at the start of the deployment phase. Note that the script will simply exit if the instance has already been enrolled in IdM. (BZ#1661635)

- In earlier releases of Red Hat OpenStack Platform, when these conditions were true:

  - The option, reclaim_instance_interval, was greater than zero.

  - The option, delete_on_termination, was set to true.

  - The instance which is booted from the volume was deleted.
    Then, after the "reclaim_instance_interval" passed, the volume on which the instance was booted, incorrectly displayed a status of "attached" and "in-use".

    In RHOSP 15, the workaround is to do the following:

1. In the Compute service configuration file, nova.conf, add a user/project configuration to the group, cinder.

2. When the context is is_admin, connect to the Block Storage service (cinder) API, authenticating with nova.conf and without using a token. (BZ#1691839)

- Because Red Hat Ceph Storage 4 is at beta when Red Hat OpenStack Platform 15 is at GA, a new configuration option has been added to RHOSP 15 to prevent any accidental deployments of Red Hat Ceph Storage 4 Beta in a production environment.
  The new Orchestration service (heat) configuration option, EnableRhcs4Beta, is set by default to "False", and therefore prevents director from deploying Red Hat Ceph Storage 4 Beta by accident. (BZ#1722036)

- When the "live_migration_wait_for_vif_plug" flag and OVN are enabled, the Red Hat OpenStack Platform Compute service (nova) times out, because the "network-vif-plugged" event never occurs.
  The workaround is to disable the "live_migration_wait_for_vif_plug" flag. Disabling this flag does not impact the live migration feature.

  When OVN is used, the default is: live_migration_wait_for_vif_plug = false. (BZ#1722041)

- In earlier Red Hat OpenStack Platform versions, when you deployed the Block Storage service (cinder) on a NetApp back end server, director warned you that deprecated parameters were specified.
  In RHOSP 15, these deprecated director parameters have been updated to align with the latest NetApp driver settings. A new parameter, CinderNetappPoolNameSearchPattern, replaces, CinderNetappStoragePools. The deprecated parameter, CinderNetappEseriesHostType, has been removed. (BZ#1595543)

- Red Hat OpenStack Platform director now has the ability to control Block Storage service (cinder) snapshots on NFS back ends. A new director parameter, CinderNfsSnapshotSupport, has a default value of True. (BZ#1633146)

- In Red Hat OpenStack Platform 15, the Image service (glance) is automatically configured for any glance-import execution to convert imported images into RAW format when Red Hat Ceph Storage is used as the backend for the Image service. (BZ#1666529)

- In Red Hat OpenStack Platform 15, you can specify MTU (maximum transmission unit) settings for each network, and RHOSP will automatically write those settings to the network interface configuration templates. MTU values should be set in the network_data.yaml file.
  This enhancement alleviates the step of manually updating the network templates for each role, and reduces the likelihood of manual entry errors. (BZ#1240852)

Changes to the puppet component:

- In Red Hat OpenStack Platform 15, director uses version 5.5 of Puppet. (BZ#1619762)

Changes to the puppet-manila component:

- The Shared File Systems service (manila) API now supports Transport Layer Security (TLS) endpoints on the internal API network, through SSL/TLS certificates. The Shared File Systems service is automatically secured when you opt to secure Red Hat OpenStack Platform during deployment. (BZ#1484601)

Changes to the puppet-nova component:

- In Red Hat OpenStack Platform 15, you are now able to customize libvirt NFS mount options for Block Storage service (cinder) volumes, using the configuration setting, nfs_mount_options. Here is an example:

  parameter_defaults: ComputeExtraConfig: nova::compute::libvirt::nfs_mount_options: "vers=4.2,lookupcache=pos" (BZ#1715094)

Changes to the puppet-tripleo component:

- In Red Hat OpenStack Platform 15, the monitoring agent, Sensu client service, is deprecated. In a future Red Hat OpenStack Platform version, the Sensu client service will be removed. (BZ#1676951)

Changes to the python-cinder-tests-tempest component:

- Before this update, Cinder consistency group tests failed because the tests used non-admin credentials. This update configures the tests to use admin credentials, allowing the consistency group tests to succeed. (BZ#1622968)

Changes to the python-networking-ovn component:

- This update fixes a bug that caused live migrations to fail.
  Before the update, with OVN enabled, a live migration could get stuck waiting for Neutron to send vif_plugged notifications.

  This update emits the vif_plugged notification under specific conditions, allowing the live migration to pass. (BZ#1743231)

Changes to the python-novajoin component:

- As a technology preview in Red Hat OpenStack Platform 15, the novajoin service tech uses the new, versioned format of notifications sent by the Compute service (nova).
  To enable the new format, set the value of the new configuration setting, configuration_format, to "versioned." The default value for configuration_format is "unversioned". * In a future version of RHOSP, unversioned notifications will be deprecated. (BZ#1624486)

- As a technology preview in Red Hat OpenStack Platform 15, the novajoin service uses the Python 3 runtime. (BZ#1624488)

Changes to the python-paunch component:

- With Paunch you can now manage container memory consumption using three new attributes: mem_limit, memswap_limit, and mem_swappiness. (BZ#1647057)

Changes to the python-tripleoclient component:

- In some earlier Red Hat OpenStack Platform versions, the following validations were not working:

  - neutron-sanity-check

  - rabbitmq-limits

  - undercloud-process-count

  - undercloud-tokenflush

  - undercloud-heat-purge-deleted

In RHOSP 15, this problem has been corrected. A new director CLInow allows you to run the earlier listed validations through Red Hat Ansible Automation directly from the Undercloud machine. (BZ#1730073)

- Because Red Hat Ceph Storage 4 is at beta when Red Hat OpenStack Platform 15 is at GA, a new configuration option has been added to RHOSP 15 to prevent any accidental deployments of Red Hat Ceph Storage 4 Beta in a production environment.
  The new Orchestration service (heat) configuration option, EnableRhcs4Beta, is set by default to "False", and therefore prevents director from deploying Red Hat Ceph Storage 4 Beta by accident. (BZ#1740715)

- Red Hat OpenStack Platform (RHOSP) does not yet support upgrading to version 15 from earlier RHOSP versions. Support for upgrading will be added to a future update of RHOSP 15. (BZ#1741244)

## 4.2. RHBA-2020:0643 — REDHAT OPENSTACK PLATFORM 15 MAINTENANCE RELEASE ADVISORY

The enhancements and bug fixes contained in this section are addressed by advisory RHBA-2020:0643. Further information about this advisory is available at link: https://access.redhat.com/errata/RHBA-2020:0643-05

**ansible-role-tripleo-modify-image**

**BZ#**1794167

There is a known issue when building container images with buildah. The default format for the image is OCI, but podman 1.6.x contains stricter restrictions about container format metadata. As a result, containers that you push to the undercloud registry can fail if they were originally in OCI format.

The workaround is to use the --format docker option to build images in docker format instead of OCI format, and you can push the containers to the undercloud registry successfully.

**diskimage-builder**

**BZ#**1775581

Previously, the nouveau kernel module was included in initramfs and conflicted with the nVidia vGPU drivers. As a result, the boot process could hang on RHOSP RHEL8 compute nodes with nVidia vGPU cards and drivers installed.

With this update, nouveau is explicitly omitted from the RHOSP initramfs.

**openstack-tripleo-heat-templates**

**BZ#**1774076

Previously, there was a change to the log parameter in the podman interface that introduced an issue with tripleo-heat-templates, which caused updates to fail.

With this update, the issue has been resolved and updates pass successfully.

**BZ#**1749443

The Compute services (nova) can fail to deploy because the nova_wait_for_compute_service script is unable to query the Nova API. If you use a remote container image registry outside of the undercloud, the Nova API service might not finish deploying in time.

The workaround is to rerun the deployment command, or to use a local container image registry on the undercloud.

BZ#1767840

Previously, changing the default membership role from Member to member caused ceph-rgw to deny access to standard users because keystone roles are case insensitive, but ceph-rgw role matching is case sensitive. As a result, users with the member role could not access ceph-rgw.

With this update, ceph-rgw accepts users with both Member and member roles.

BZ#1746273

Previously, deploying the stack with all networks disabled failed because the 'cloud_name_{{network.name_lower}}' property was defined for disabled networks.

With this update, the 'cloud_name_{{network.name_lower}}' property is no longer added for disabled networks and deployments are successful.

BZ#1751809

With this update, the credentials that you supply in the ContainerImageRegistryCredentials parameter pass to ceph-ansible automatically if the registry name matches the registry name in the ceph_namespace parameter.

BZ#1765438

In Red Hat OpenStack Platform (RHOSP) 15, the rhel-registration pre-deployment script has been removed because it is not compatible with RHEL8. Use the Ansible-based overcloud registration method instead:

https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/15/html/advanced_overcloud_customization/ansible-based-registration

**openstack-tripleo-image-elements**

BZ#1759430

Previously, the default NIC naming changed to use non-persistent node names such as ethXX instead of consistent names like enoX, ensXfY, ensX. As a result, the NIC names present in the introspection data did not match the overcloud NIC names.

With this update, the setting net.ifnames=0 has been removed from grub-config in the overcloud image and the introspection data contains consistent NIC names.

**openstack-tripleo-validations**

BZ#1764260

> Previously, the default values for dhcp_start and dhcp_end in the undercloud.conf file did not provide enough IP addresses to pass tripleo validation and the ctlplane-ip-range validation failed.
>
> With this update, the IP range is larger and validation passes successfully.

**os-net-config**

BZ#1718108

> Previously, routes on SR-IOV PF interfaces were not set properly and these routes were ignored.
>
> With this update, routes for SR-IOV PF interfaces function correctly.