Red Hat OpenStack Platform 13

Release Notes

Release details for Red Hat OpenStack Platform 13

Last Updated: 2018-09-20

Red Hat OpenStack Platform 13 Release Notes

Release details for Red Hat OpenStack Platform 13

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 "Queens" 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 "Queens" release itself are available at the following location: https://releases.openstack.org/queens/index.html.

Red Hat OpenStack Platform uses components from other Red Hat products. See the following links 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 URL for more details on the add-on: http://www.redhat.com/products/enterprise-linux-add-ons/high-availability/. See the following URL for details on the package versions to use in combination with Red Hat OpenStack Platform: https://access.redhat.com/site/solutions/509783

1.2. REQUIREMENTS

Red Hat OpenStack Platform supports the most recent release of Red Hat Enterprise Linux. This version of Red Hat OpenStack Platform is supported on Red Hat Enterprise Linux 7.5.

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
- Firefox
- 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. BARE METAL PROVISIONING SUPPORTED OPERATING SYSTEMS

For a list of the supported 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).

1.8. HYPERVISOR SUPPORT

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

Ironic has been fully supported since the release of Red Hat OpenStack Platform 7 (Kilo). Ironic allows you to provision bare-metal machines using common technologies (such as PXE boot 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, and non-KVM libvirt hypervisors.

1.9. CONTENT DELIVERY NETWORK (CDN) REPOSITORIES

This section describes the repository settings required to deploy Red Hat OpenStack Platform 13.

You can install Red Hat OpenStack Platform 13 through the Content Delivery Network (CDN). To do so, configure **`subscription-manager`** to use the correct repositories.

Run the following command to enable a CDN repository:

```
#subscription-manager repos --enable=[reponame]
```

Run the following command to disable a CDN repository:

```
#subscription-manager repos --disable=[reponame]
```

Table 1.1. Required Repositories (x86_64)

Repository Name	Repository Label
Red Hat Enterprise Linux 7 Server (RPMS)	`rhel-7-server-rpms`
Red Hat Enterprise Linux 7 Server - RH Common (RPMs)	`rhel-7-server-rh-common-rpms`
Red Hat Enterprise Linux High Availability (for RHEL 7 Server)	`rhel-ha-for-rhel-7-server-rpms`
Red Hat OpenStack Platform 13 for RHEL 7 (RPMs)	`rhel-7-server-openstack-13-rpms`
Red Hat Enterprise Linux 7 Server - Extras (RPMs)	`rhel-7-server-extras-rpms`

Table 1.2. Optional Repositories (x86_64)

Repository Name	Repository Label
Red Hat Enterprise Linux 7 Server - Optional	`rhel-7-server-optional-rpms`
Red Hat OpenStack Platform 13 Operational Tools for RHEL 7 (RPMs)	`rhel-7-server-openstack-13-optools-rpms`

Table 1.3. Required Repositories (ppc64le)

Repository Name	Repository Label
Red Hat Enterprise Linux for IBM Power, little endian	`rhel-7-for-power-le-rpms`
Red Hat OpenStack Platform 13 for RHEL 7 (RPMs)	`rhel-7-server-openstack-13-for-power-le-rpms`

Repositories to Disable

The following table outlines the repositories you must disable to ensure Red Hat OpenStack Platform 13 functions correctly.

Table 1.4. Repositories to Disable

Repository Name	Repository Label
Red Hat CloudForms Management Engine	`"cf-me-*"`
Red Hat Enterprise Virtualization	`"rhel-7-server-rhev*"`

Repository Name	Repository Label
Red Hat Enterprise Linux 7 Server - Extended Update Support	`"*-eus-rpms"`

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. 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 OpenStack deployment. Facilities available via the Customer Portal include:

- Knowledge base articles and solutions.
- Technical briefs.
- Product documentation.
- 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 OpenStack 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 OPENSTACK PLATFORM DIRECTOR

This section outlines the top new features for the director.

Fast forward upgrades

The director provides a **fast forward upgrade** path through multiple versions, specifically from **Red Hat OpenStack Platform 10** to **Red Hat OpenStack Platform 13**. The goal is to provide users an opportunity to remain on certain OpenStack versions that are considered **long life versions** and upgrade when the next long life version is available. Full instructions are available in the *Fast Forward Upgrades* Guide.

Red Hat Virtualization control plane

The director now supports provisioning an overcloud using Controller nodes deployed in Red Hat Virtualization. For more information about new virtualization features, see Virtualize your OpenStack control plane with Red Hat Virtualization and Red Hat OpenStack Platform 13.

2.2. CONTAINERS

This section outlines the top new features for containerization in Red Hat OpenStack Platform.

Fully containerized services

The release provides all Red Hat OpenStack Platform services as containers, including services that were not containerized in the previous version: OpenStack Networking (neutron), OpenStack Block Storage (cinder), and OpenStack Shared File Systems (manila). The overcloud now uses fully containerized services.

2.3. BARE METAL SERVICE

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

L3 routed spine-leaf network

The director includes the capability to define multiple networks for provisioning and introspection functions. This feature, in conjunction with composable networks, allows users to provision and configure a complete L3 routed spine-leaf architecture for the overcloud. Full instructions are available in the *Spine Leaf Networking* Guide.

Red Hat Virtualization driver

The director OpenStack Bare Metal (ironic) service includes a driver (`staging-ovirt`) to manage virtual nodes within a Red Hat Virtualization environment.

2.4. CEPH STORAGE

This section outlines the top new features for Ceph Storage.

Red Hat Ceph Storage 3.0 support

With this release, Red Hat Ceph Storage 3.0 (luminous) is the default supported version of Ceph for Red Hat OpenStack and is the default version deployed by director. Ceph now supports rolling upgrades from version 2.x to 3. Upgrading to the new OpenStack release also upgrades Red Hat Ceph Storage to 3.0 if your Ceph cluster was deployed using director.

Scale out Ceph Metadata Server and RADOS Gateway nodes

Red Hat Ceph Storage 3.0 adds support for scaling metadata load across multiple metadata servers (MDS) by appropriate configuration of the Ceph File System (CephFS). Once configured, extra dedicated MDS servers available in your Ceph cluster are automatically assigned to take on this extra load. Additionally, new dedicated Ceph RADOS Gateway (RGW) nodes can be added, allowing RGW to scale up as needed.

Manila CephFS storage with NFS

The Shared File System service (manila) supports mounting shared file systems backed by a Ceph File System (CephFS) via the NFSv4 protocol. NFS-Ganesha servers operating on Controller nodes are used to export CephFS to tenants with High Availability (HA). Tenants are isolated from one another and may only access CephFS through the provided NFS gateway interface. This new feature is fully integrated into director, thereby enabling CephFS back end deployment and configuration for the Shared File System service.

Enhanced multiple Cinder Ceph pools support

Block Storage (cinder) RADOS block device (RBD) back ends can be mapped to different pools within the same Ceph cluster using a director template parameter, **`CinderRbdExtraPools`**. A new Block Storage RBD back end is created for each Ceph pool associated with this parameter, in addition to the standard RBD back end associated with the **`CinderRbdPoolName`** parameter.

RBD mirror director with ceph-ansible

The Ceph **rbd-mirror** daemon pulls image updates from a remote cluster and applies them to the image within a local cluster. RBD mirror is deployed as a container using **ceph-ansible** with Red Hat Ceph Storage 3.0 (luminous). OpenStack metadata related to the image is not copied by **rbd-mirror**.

2.5. COMPUTE

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

Real-Time KVM integration

Integration of real time KVM (RT-KVM) with the Compute service is now fully supported. RT-KVM benefits are:

- Deterministic and low average latency for system calls and interrupts.
- Precision Time Protocol (PTP) support in the guest instance for accurate clock synchronization (community support for this release).

2.6. HIGH AVAILABILITY

This section outlines the top new features for high availability.

Director integration for Instance HA

You can now deploy Instance HA with the director. This allows you to configure installation and upgrade for Instance HA without further manual steps.

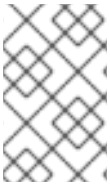

NOTE

Director integration for Instance HA is available only from version 13 and later. To upgrade from previous versions to version 13, including fast-forward upgrades, you must first manually disable Instance HA.

2.7. METRICS AND MONITORING

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

`collectd` 5.8 integration

The **`collectd`** 5.8 version includes the following additional plugins:

- **`ovs-stats`** - The plugin collects the statistics of OVS connected bridges and interfaces.
- **`ovs-events`** - The plugin monitors the link status of Open vSwitch (OVS) connected interfaces, dispatches the values to **`collectd`**, and sends the notification whenever the link state change occurs in the OVS database.
- **`hugepages`** - The **`hugepages`** plugin allows the monitoring of free and used hugepages by numbers, bytes, or percentage on a platform.
- **`intel-rdt`** - The **`intel_rdt`** plugin collects information provided by monitoring features of Intel Resource Director Technology (Intel® RDT) like Cache Monitoring Technology (CMT), Memory Bandwidth Monitoring (MBM). These features provide information about shared resource usage such as last level cache occupancy, local memory bandwidth usage, remote memory bandwidth usage, and instructions per clock.
- **`libvirt`** plugin extension - The **`libvirt`** plugin is extended to support CMT, MBM, CPU Pinning, Utilization, and State metrics on the platform.

`collectd` and `gnocchi` integration

The **`collectd-gnocchi`** plugin sends the metrics to gnocchi. By default, it creates a resource type named ***`collectd`*** and a new resource for each host monitored.
Each host has a list of metrics created dynamically using the following naming convention:

```
plugin-plugin_instance/type-type_instance-value_number
```

In order for the metrics to be created properly, make sure the archive policy rules match.

Support `sensu` with multiple RabbitMQ servers

With this release, the Red Hat OpenStack Platform adds support to **`sensu`** with multiple RabbitMQ servers. To achieve this, you need to use the **`MonitoringRabbitCluster`** parameter in the **`config.yaml`** file.

Intel Resource Director Technology/Memory Bandwidth Monitoring support

Memory Bandwidth Monitoring (MBM) is an integral part of the Intel® Resource Director Technology (RDT). Memory usage and availability is gathered from all the nodes and made available to OpenStack to make better scheduling decisions and deliver on SLAs.

2.8. NETWORK FUNCTIONS VIRTUALIZATION

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

Real-Time KVM Compute role for NFV workloads

The real-time KVM (RT-KVM) Compute nodes now support NFV workloads, with the addition of a RT-KVM Compute node role. This new role exposes a subset of Compute nodes with real-time capabilities to support guests with stringent latency requirements.

2.9. OPENDAYLIGHT

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

OpenDaylight integration

OpenDaylight is a flexible, modular, and open SDN platform, that is now fully supported with this Red Hat OpenStack Platform release. The current Red Hat offering combines carefully selected OpenDaylight components that are designed to enable the OpenDaylight SDN controller as a networking backend for OpenStack. The key OpenDaylight project used in this solution is NetVirt, with support for the OpenStack neutron API.
For more information, see the Red Hat OpenDaylight Product Guide and the Red Hat OpenDaylight Installation and Configuration Guide.

2.10. OPENSTACK NETWORKING

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

Octavia LBaaS

Octavia is now fully supported. Octavia is an official OpenStack project that provides load balancing capabilities and is intended to replace the current HAProxy-based implementation. Octavia implements the LBaaS v2 API, but also provides additional features. Octavia includes a reference load balancing driver that provides load balancing with *amphora* (implemented as Compute VMs).

Open Virtual Network (OVN)

OVN is now fully supported. OVN is an Open vSwitch-based network virtualization solution for supplying network services to instances. OVN fully supports the **`neutron`** API.

2.11. SECURITY

This section outlines the top new features for security components.

Barbican

OpenStack Key Manager (barbican) is a secrets manager for Red Hat OpenStack Platform. You can use the barbican API and command line to centrally manage the certificates, keys, and passwords used by OpenStack services.

Barbican - Support for encrypted volumes

You can use barbican to manage your Block Storage (cinder) encryption keys. This configuration uses LUKS to encrypt the disks attached to your instances, including boot disks. The key management aspect is performed transparently to the user.

Barbican - glance image signing

You can configure the Image Service (glance) to verify that an uploaded image has not been tampered with. The image is first signed with a key that is stored in barbican, with the image then being validated before each use.

Integration with Policy Decision Points (PDP)

For customers that rely on Policy Decision Points (PDP) to control access to resources, Identity Service (keystone) can now integrate projects with an external PDP for authorization checks. The external PDP can evaluate access requests and can grant or deny access based on established policy.

Infrastructure and virtualization hardening

AIDE Intrusion detection is now available under tech preview. The director's AIDE service allows an operator to centrally set their intrusion detection ruleset and then install and setup AIDE on the overcloud.

2.12. STORAGE

This section outlines the top new features for storage components.

Block Storage - Containerized deployment of the Block Storage service

Containerized deployment of the Block Storage service (cinder) is now the default in this release. If you use a back end for these services that has external installation dependencies, you must obtain vendor-specific containers for your deployment.

Block Storage - Multi-back end availability zones

The Block Storage service (cinder) now allows back end availability zones to be defined using a new driver configuration option, **`backend_availability_zone`**, in the back end sections of the configuration file. In previous versions, back ends configured in a cinder-volume had to be part of the same storage availability zone.

Block Storage - OpenStack Key Manager support

The Block Storage service (cinder) can now use the OpenStack Key Manager (barbican) to store encryption keys used for volume encryption. This feature is enabled by configuring the OpenStack Key Manager in director. New keys can be added to the OpenStack Key Manager by users with the admin or creator roles by Identity Service (keystone).

Block Storage - RBD driver encryption support

The RBD driver now handles Block Storage service (cinder) volume encryption using LUKS. This feature provides the capability to encrypt volumes on RBD using the Block Storage service and Compute service, providing data-at-rest security. The OpenStack Key Manager (barbican) is required to use RBD driver encryption. RBD driver encryption is only supported for the Block Storage service.

Image Service - Image signing and verification support

The Image Service (glance) now provides signing and signature validation of bootable images using OpenStack Key Manager (barbican). Image signatures are now verified prior to storing the image. You must add an encryption signature to the original image before uploading it to the Image Service. This signature is used to validate the image upon booting. OpenStack Key Manager provides key management support for signing keys.

Object Storage - At-rest encryption and OpenStack Key Manager support

The Object Storage (swift) service can now store objects in encrypted form using AES in CTR mode with 256-bit keys stored in the OpenStack Key Manager (barbican). Once encryption is enabled for Object Storage using director, the system creates a single key used to encrypt all objects in the cluster. This provides options for protecting objects and maintaining security compliance in Object Storage clusters.

Shared File System - Containerized deployment of the Shared File System service

Containerized deployment of the Shared File System service (manila) is now the default in this release. If you use a back end for these services that has external installation dependencies, you must obtain vendor-specific containers for your deployment.

Shared File System - IPv6 access rule support with NetApp ONTAP cDOT driver

The Shared File System service (manila) now supports exporting shares backed by NetApp ONTAP back ends over IPv6 networks. Access to the exported shares is controlled by IPv6 client addresses.

Shared File System - Manila CephFS storage with NFS

The Shared File System service (manila) supports mounting shared file systems backed by a Ceph File System (CephFS) via the NFSv4 protocol. NFS-Ganesha servers operating on Controller nodes are used to export CephFS to tenants with High Availability (HA). Tenants are isolated from one

another and may only access CephFS through the provided NFS gateway interface. This new feature is fully integrated into director, thereby enabling CephFS back end deployment and configuration for the Shared File System service.

2.13. TECHNOLOGY PREVIEWS

This section outlines features that are in technology preview in Red Hat OpenStack Platform 13.

NOTE

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

2.13.1. New Technology Previews

The following new features are provided as technology previews:

Ansible-based configuration (config download)

The director can now generate a set of Ansible playbooks using an overcloud plan as a basis. This changes the overcloud configuration method from OpenStack Orchestration (heat) to an Ansible-based method. Some supported OpenStack Platform 13 features, such as upgrades, use this feature as part of their processes. However, usage outside of these supported areas is not recommended for production and only available as a technology preview.

OVS hardware offload

Open vSwitch (OVS) hardware offload accelerates OVS by moving heavy processing to hardware with SmartNICs. This saves host resources by offloading the OVS processing to the SmartNIC.

2.13.2. Previously Released Technology Previews

The following features remain as technology previews:

Benchmarking service

Rally is a benchmarking tool that automates and unifies multi-node OpenStack deployment, cloud verification, benchmarking, and profiling. It can be used as a basic tool for an OpenStack CI/CD system that would continuously improve its SLA, performance, and stability. It consists of the following core components:

- Server Providers - provide a unified interface for interaction with different virtualization technologies (LXS, Virsh etc.) and cloud suppliers. It does so via ssh access and in one L3 network.
- Deploy Engines - deploy an OpenStack distribution before any benchmarking procedures take place, using servers retrieved from Server Providers.
- Verification - runs specific set of tests against the deployed cloud to check that it works correctly, collects results and presents them in human readable form.
- Benchmark Engine - allows you to write parameterized benchmark scenarios and run them against the cloud.

Benchmarking service - introduction of a new plug-in type: hooks

Allows test scenarios to run as iterations, and provides timestamps (and other information) about executed actions in the rally report.

Benchmarking service - new scenarios

Benchmarking scenarios have been added for nova, cinder, magnum, ceilometer, manila, and neutron.

Benchmarking service - refactor of the verification component

Rally Verify is used to launch Tempest. It was refactored to cover a new model: verifier type, verifier, and verification results.

Cells

OpenStack Compute includes the concept of Cells, provided by the **`nova-cells`** package, for dividing computing resources. In this release, Cells v1 has been replaced by Cells v2. Red Hat OpenStack Platform deploys a "cell of one" as a default configuration, but does not support multi-cell deployments at this time.

DNS-as-a-Service (DNSaaS)

DNS-as-a-Service (DNSaaS), also known as Designate, includes a REST API for domain and record management, is multi-tenanted, and integrates with OpenStack Identity Service (keystone) for authentication. DNSaaS includes a framework for integration with Compute (nova) and OpenStack Networking (neutron) notifications, allowing auto-generated DNS records. DNSaaS includes integration with the Bind9 back end.

Firewall-as-a-Service (FWaaS)

The Firewall-as-a-Service plug-in adds perimeter firewall management to OpenStack Networking (neutron). FWaaS uses iptables to apply firewall policy to all virtual routers within a project and supports one firewall policy and logical firewall instance per project. FWaaS operates at the perimeter by filtering traffic at the OpenStack Networking (neutron) router. This distinguishes it from security groups, which operate at the instance level.

Google Cloud storage backup driver (Block Storage)

The Block Storage (cinder) service can now be configured to use Google Cloud Storage for storing volume backups. This feature presents an alternative to the costly maintenance of a secondary cloud simply for disaster recovery.

Link aggregation for bare metal nodes

This release introduces link aggregation for bare metal nodes. Link aggregation allows you to configure bonding on your bare metal node NICs to support failover and load balancing. This feature requires specific hardware switch vendor support that can be configured from a dedicated neutron plug-in. Verify that your hardware vendor switch supports the correct neutron plug-in.
Alternatively, you can manually preconfigure switches to have bonds set up for the bare metal nodes. To enable nodes to boot off one of the bond interfaces, the switches need to support both LACP and LACP fallback (bond links fall back to individual links if a bond is not formed). Otherwise, the nodes will also need a separate provisioning and cleaning network.

Red Hat OpenStack Platform for POWER

You can now deploy pre-provisioned overcloud Compute nodes on IBM POWER8 little endian hardware.

Red Hat SSO

This release includes a version of the keycloak-httpd-client-install package. This package provides a command-line tool that helps configure the Apache mod_auth_mellon SAML Service Provider as a client of the Keycloak SAML IdP.

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 13 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#1419556

```
The Object Store service (swift) can now integrate with Barbican to
transparently encrypt and decrypt your stored (at-rest) objects. At-rest
encryption is distinct from in-transit encryption and refers to the
objects being encrypted while being stored on disk.

Swift objects are stored as clear text on disk. These disks can pose a
security risk if not properly disposed of when they reach end-of-life.
Encrypting the objects mitigates that risk.

Swift performs these encryption tasks transparently, with the objects
being automatically encrypted when uploaded to swift, then automatically
decrypted when served to a user. This encryption and decryption is done
using the same (symmetric) key, which is stored in Barbican.
```

BZ#1540239

```
This enhancement adds support for sending metrics data to a Gnocchi DB
instance.

The following new parameters for collectd composable service were added.
If CollectdGnocchiAuthMode is set to 'simple', then
CollectdGnocchiProtocol, CollectdGnocchiServer, CollectdGnocchiPort and
CollectdGnocchiUser are taken into account for configuration.

If CollectdGnocchiAuthMode is set to 'keystone', then
CollectdGnocchiKeystone* parameters are taken into account for
configuration.

Following is a detailed description of added parameters:

  CollectdGnocchiAuthMode:
    type: string
    description: >
```

```
      Type of authentication Gnocchi server is using. Supported values are
      'simple' and 'keystone'.
    default: 'simple'
  CollectdGnocchiProtocol:
    type: string
    description: API protocol Gnocchi server is using.
    default: 'http'
  CollectdGnocchiServer:
    type: string
    description: >
      The name or address of a gnocchi endpoint to which we should
      send metrics.
    default: nil
  CollectdGnocchiPort:
    type: number
    description: The port to which we will connect on the Gnocchi server.
    default: 8041
  CollectdGnocchiUser:
    type: string
    description: >
      Username for authenticating to the remote Gnocchi server using
simple
      authentication.
    default: nil
  CollectdGnocchiKeystoneAuthUrl:
    type: string
    description: Keystone endpoint URL to authenticate to.
    default: nil
  CollectdGnocchiKeystoneUserName:
    type: string
    description: Username for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneUserId:
    type: string
    description: User ID for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystonePassword:
    type: string
    description: Password for authenticating to Keystone
    default: nil
  CollectdGnocchiKeystoneProjectId:
    type: string
    description: Project ID for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneProjectName:
    type: string
    description: Project name for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneUserDomainId:
    type: string
    description: User domain ID for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneUserDomainName:
    type: string
    description: User domain name for authenticating to Keystone.
    default: nil
```

```
  CollectdGnocchiKeystoneProjectDomainId:
    type: string
    description: Project domain ID for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneProjectDomainName:
    type: string
    description: Project domain name for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneRegionName:
    type: string
    description: Region name for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneInterface:
    type: string
    description: Type of Keystone endpoint to authenticate to.
    default: nil
  CollectdGnocchiKeystoneEndpoint:
    type: string
    description: >
      Explicitly state Gnocchi server URL if you want to override
      Keystone value
    default: nil
  CollectdGnocchiResourceType:
    type: string
    description: >
      Default resource type created by the collectd-gnocchi plugin in
Gnocchi
      to store hosts.
    default: 'collectd'
  CollectdGnocchiBatchSize:
    type: number
    description: Minimum number of values Gnocchi should batch.
    default: 10
```

BZ#1592823

```
Logs from Ansible playbooks now include timestamps that provide
information about the timing of actions during deployment, updates, and
upgrades.
```

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#1446311

```
This release adds support for PCI device NUMA affinity policies, which are
configured as part of the “[pci]alias” configuration options. Three
policies are supported:

“required” (must have)
“legacy” (default; must have, if available)
“preferred” (nice to have)
```

```
In all cases, strict NUMA affinity is provided, if possible. These
policies allow you to configure how strict your NUMA affinity should be
per PCI alias to maximize resource utilization. The key difference between
the policies is how much NUMA affinity you're willing to forsake before
failing to schedule.

When the “preferred” policy is configured for a PCI device, nova uses CPUs
on a different NUMA node from the NUMA node of the PCI device, if it is
available. This results in increased resource utilization, but performance
is reduced for these instances.
```

BZ#1488095

```
From RHOS-12 onwards, the OpenStack services are becoming containerized. In
this release, we containerize OpenStack Tempest as well. The containerized
OpenStack Tempest is available as a Technology Preview.
```

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#1468020

```
The Shared File System service (manila) now provides IPv6 access rule
support with NetApp ONTAP cDOT driver, which lets you use manila with IPv6
environments.

As a result, the Shared File System service now supports exporting shares
backed by NetApp ONTAP back ends over IPv6 networks. Access to the
exported shares is controlled by IPv6 client addresses.
```

BZ#1469208

```
The Shared File System service (manila) supports mounting shared file
systems backed by a Ceph File System (CephFS) via the NFSv4 protocol. NFS-
Ganesha servers operating on Controller nodes are used to export CephFS to
tenants with high availability (HA). Tenants are isolated from one another
and may only access CephFS through the provided NFS gateway interface.
This new feature is fully integrated into director, enabling CephFS back
end deployment and configuration for the Shared File System service.
```

BZ#1496584

```
When neutron services are containerized, trying to run commands in a
network namespace might fail with t he following error:

# ip netns exec qrouter...
RTNETLINK answers: Invalid argument

In order to run a command inside a network namespace, you must do it from
the neutron container that created the namespace.  For example, the l3-
```

> agent creates network namespace for routers, so the command would need to change to:
>
> # docker exec neutron_l3_agent ip netns exec qrouter...
>
> Similarly with network namespaces beginning with 'qdhcp' you would need to exec from the 'neutron_dhcp' container.

BZ#1503521

> This version introduces support for internal DNS resolution in networking-ovn. Although there are two know limitations,
> one is bz#1581332 which prevents proper resolution of internal fqdn requests via internal dns.
>
> Please note that the extension is not configured by default by tripleo on the GA release. See bz#1577592 for a workaround.

BZ#1533206

> The openstack-gnocchi packages have been renamed to gnocchi. The openstack- prefix was removed because of an upstream project scoping change. Gnocchi has been moved out of the OpenStack umbrella and is maintained as a stand-alone project.

BZ#1556933

> Since version 2.1, python-cryptography checks that the CNS Names used in certificates are compliant with IDN standards. If the found names do not follow this specification, cryptography will fail to validate the certificate and different errors may be found when using OpenStack command line interface or in OpenStack service logs.

BZ#1563412

> The reserved host memory for OpenStack Compute (nova) has increased from 2048 MB to 4096 MB. This can affect capacity estimations for your environment. If necessary, you can reconfigure the reserved memory using the 'NovaReservedHostMemory' parameter in a environment file. For example:
>
> ```
> parameter_defaults:
>   NovaReservedHostMemory: 2048
> ```

BZ#1564176

> The python-mistralclient is not part of any supported overcloud use-cases so it is being dropped from the -tools channels for the OSP 13 release.

BZ#1567735

> OSP13 using OVN as the networking backend won't include IPv6 support in the first release. There is a problem with the responses to the Neighbor Solicitation requests coming from guests VMs which causes a loss of the default routes.

-

BZ#1575752

> In previous versions, the *NetName parameters (e.g. InternalApiNetName) changed the names of the default networks. This is no longer supported.
>
> To change the names of the default networks, use a custom composable network file (network_data.yaml) and include it with your 'openstack overcloud deploy' command using the '-n' option. In this file you should set the "name_lower" field to the custom net name for the network you want to change. For more information, see "Using Composable Networks" in the Advanced Overcloud Customization guide.
>
> In addition, you need to add a local parameter for the ServiceNetMap table to network_environment.yaml and override all the default values for the old network name to the new custom name. The default values can be found in /usr/share/openstack-tripleo-heat-templates/network/service_net_map.j2.yaml. This requirement to modify ServiceNetMap will not be necessary in future OSP-13 releases.

BZ#1577537

> Fixes OSP 13 Beta issue where some container images were not available.

BZ#1578312

> When the OVSDB server fails over to a different controller node, a reconnection from neutron-server/metadata-agent does not take place because they are not detecting this condition.
>
> As a result, booting VMs may not work as metadata-agent will not provision new metadata namespaces and the clustering is not behaving as expected.
>
> A possible workaround is to restart the ovn_metadata_agent container in all the compute nodes after a new controller has been promoted as master for OVN databases. Also increase the ovsdb_probe_interval on the plugin.ini to a value of 600000 milliseconds.

BZ#1589849

> When the OVN metadata agent is stopped in a Compute node, all the VMs on that node will not have access to the metadata service. The impact is that if a new VM is spawned or an existing VM is rebooted, the VM will fail to access metadata until the OVN metadata agent is brought up back again.

BZ#1592528

> In rare circumstances, after rebooting controller nodes several times, RabbitMQ may be running in an inconsistent state that will block API operations on the overcloud.
>
> The symptoms for this issue are:
> - Entries in any of the OpenStack service logs of the form:
> DuplicateMessageError: Found duplicate

```
message(629ff0024219488499b0fac0cacaa3a5). Skipping it.
 - "openstack network agent list" returns that some agents are DOWN

To restore normal operation, run the following command on any of the
controller nodes (you only need to do this on one controller):
 pcs resource restart rabbitmq-bundle
```

3.1.4. Known Issues

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

BZ#1321179

```
OpenStack command-line clients that use `python-requests` can not currently
validate certificates that have an IP address in the SAN field.
```

BZ#1461132

```
When using Red Hat Ceph Storage as a Block Storage backend for both Cinder
volume and Cinder backup, any attempts to perform an incremental backup
will result in a
full backup instead, without any warning. This is a known issue.
```

BZ#1508449

```
OVN serves DHCP as an openflow controller with ovn-controller directly on
compute nodes. But SR-IOV instances are directly attached to the network
through the VF/PF. As such, SR-IOV instances will not be able to get DHCP
responses from anywhere.

To workaround this issue, change  OS::TripleO::Services::NeutronDhcpAgent
to:

   OS::TripleO::Services::NeutronDhcpAgent: docker/services/neutron-
dhcp.yaml
```

BZ#1515815

```
When the router gateway is cleared, the Layer 3 flows related to learned
IP addresses is not removed. The learned IP addresses include the PNF and
external gateway IP addresses. This leads stale flows, but not any
functional issue. The external gateway and IP address does not change
frequently. The stale flows will be removed when the external network is
deleted.
```

BZ#1518126

```
Redis is unable to correctly replicate data across nodes in a HA
deployment with TLS enabled. Redis follower nodes will not contain any
data from the leader node. It is recommended to disable TLS for Redis
deployments.
```

BZ#1519783

-

> Neutron may issue an error claiming that the Quota has been exceed for Neutron Router creation. This is a known issue where multiple router resources are created with a single create request in Neutron DB due to a bug with networking-odl. The workaround for this issue is to delete the duplicated routers using the OpenStack Neutron CLI and create a router again, resulting with a single instance.

BZ#1557794

> A regression was identified in the procedure for backing up and restoring the director undercloud. As a result, the procedure requires modification and verification before it can be published.
>
> The book 'Back Up and Restore the Director Undercloud' is therefore not available with the general availability of Red Hat OpenStack Platform 13. The procedure will be updated as a priority after the general availability release, and published as soon as it is verified.

BZ#1559055

> OpenDaylight logging might be missing earlier logs. This is a known issue with journald logging of OpenDaylight (using the “docker logs opendaylght_api” command). The current workaround is to switch OpenDaylight logging to the “file” mechanism which will log inside of the container to /opt/opendaylight/data/logs/karaf.log. To do this, configure the following heat parameter: OpenDaylightLogMechanism: ‘file’.

BZ#1568012

> Connecting to an external IP fails when associating a floating IP to an instance then disassociating the floating IP. This situation happens in a tenant VLAN network when:
> * a VM spawned on a non-NAPT switch is associated with a floating IP and
> * the floating IP is removed.
> This results in a missing flow (sporadically) in the FIB table of NAPT switch.
>
> Due to the missing FIB table entry, the VM loses connectivity to the public network.
>
> Associating the floating IP to the VM restores connectivity to the public network. As long as the floating IP is associated with the VM, it will be able to connect to the internet. However, you will lose a public IP/floating IP from the external network.

BZ#1568311

> Layer 3 connectivity between nova instances across multiple subnets may fail when an instance without a floating IP tries to reach another instance that has a floating IP on another router. This occurs when nova instances are spread across multiple compute nodes. There is no suitable workaround for this issue.

BZ#1568976

> During deployment, one or more OpenDaylight instances may fail to start correctly due to a feature loading bug. This may lead to a deployment or functional failure.
>
> When a deployment passes, only two of the three OpenDaylight instances must be functional for the deployment to succeed. It is possible that the third OpenDaylight instance started incorrectly. Check the health status of each container with the `docker ps` command. If it is unhealthy, restart the container with `docker restart opendaylight_api`.
>
> When a deployment fails, the only option is to restart the deployment. For TLS-based deployments, all OpenDaylight instances must boot correctly or deployment will fail.

BZ#1571864

> Temporary removal of Heat stack resources during fast-forward upgrade preparation triggers RHEL unregistration.
> As a result, RHEL unregistration is stalled because Heat software deployment signalling does not work properly.
>
> To avoid the problem, while the overcloud is still on OSP 10 and ready to perform the last overcloud minor version update:
> 1. Edit the template file /usr/share/openstack-tripleo-heat-templates/extraconfig/pre_deploy/rhel-registration/rhel-registration.yaml
> 2. Delete RHELUnregistration and RHELUnregistrationDeployment resources from the template.
> 3. Proceed with the minor update and fast-forward upgrade procedure.

BZ#1573597

> A poorly performing Swift cluster used as a Gnocchi back end can generate 503 errors in the collectd log and "ConnectionError: ('Connection aborted.', CannotSendRequest())" errors in in gnocchi-metricd.conf.
> To mitigate the problem, increase the value of the CollectdDefaultPollingInterval parameter or improve the Swift cluster performance.

BZ#1574708

> When an OpenDaylight instance is removed from a cluster and reconnected, the instance may not successfully join the cluster. The node will eventually re-join the cluster.
>
> The following actions should be taken in such a situation:
> * Restart the faulty node.
> * Monitor the REST endpoint to verify the cluster member is healthy: http://$ODL_IP:8081/jolokia/read/org.opendaylight.controller:Category=ShardManager,name=shard-manager-config,type=DistributedConfigDatastore
> * The response should contain a field “SyncStatus”, and a value of “true” will indicate a healthy cluster member.

BZ#1574725

```
When multiple VMs in the same subnet of a VLAN provider network are
scheduled to two different Compute nodes, ARP between the VMs fails
sporadically.

Since ARP packets between those VMs fails, there is essentially no
networking between the two VMs.
```

BZ#1575023

```
The manila-share service fails to initialize because changes to ceph-
ansible's complex ceph-keys processing generate incorrect content in the
/etc/ceph/ceph.client.manila.keyring file.

To allow the manila-share service to initialize:
1) Make a copy of /usr/share/openstack/tripleo-heat-templates to use for
the overcloud deploy.

2) Edit the .../tripleo-heat-templates/docker/services/ceph-ansible/ceph-
base.yaml file to change all triple backslashes in line 295 to single
backslashes.
Before:
mon_cap: 'allow r, allow command \\\"auth del\\\", allow command \\\"auth
caps\\\", allow command \\\"auth get\\\", allow command \\\"auth get-or-
create\\\"'
After:
mon_cap: 'allow r, allow command \"auth del\", allow command \"auth
caps\", allow command \"auth get\", allow command \"auth get-or-create\"'

3) Deploy the overcloud substituting the path to the copy of tripleo-heat-
templates wherever /usr/share/openstack-tripleo-heat templates occurred in
your original overcloud-deploy command.

The ceph key /etc/ceph/ceph.client.manila.keyring file will have proper
contents and the manila-share service will initialize properly.
```

BZ#1575118

```
Ceph Release 12.2.1 lowers the maximum number of PGs allowed for each OSD.
The lower limit may cause the monitor to prematurely issue a HEALTH_WARN
message.

The monitor warning threshold has been reduced from 300 to 200 PGs per
OSD. 200 is still twice the generally recommended target of 100 PGs per
OSD. This limit can be adjusted via the mon_max_pg_per_osd option on the
monitors. The older mon_pg_warn_max_per_osd option has been removed.

The amount of PGs consumed by a pool can not be decreased. If the upgrade
causes a pre-existing deployment to reach the maximum limit, you can raise
the limit to its pre-upgrade value during the ceph-upgrade step. In an
environment file, add a parameter setting like this:

  parameter_defaults:
    CephConfigOverrides:
      mon_max_pg_per_osd: 300
```

> The setting is applied into ceph.conf and the cluster stays in HEALTH_OK state.

BZ#1575150

> There is a known issue where the OpenDaylight cluster may stop responding for up to 30 minutes when an OpenDaylight cluster member is stopped (due to failure or otherwise). The workaround is wait until the cluster becomes active again.

BZ#1575496

> When using a physical host interface for external network with Director, if the interface is not attached to an OVS bridge, the interface will not pass traffic in an OpenDaylight setup. Traffic will not pass and you should avoid this type of configuration.
>
> Always use an OVS bridge in the NIC templates for an overcloud external network. This bridge is named "br-ex" by default in Director (although you may use any name). You should attach the physical host interface used for the external network to this OVS bridge.
>
> When you use an interface attached to an OVS bridge, the deployment will function correctly and the external network traffic to tenants will work correctly.

BZ#1577975

> OpenDaylight may experience periods of very high CPU usage. This issue should not affect the functionality of OpenDaylight, although it could potentially impact other system services.

BZ#1579025

> OVN pacemaker Resource Agent (RA) script sometimes does not handle the promotion action properly when pacemaker tries to promote a slave node. This is seen when the ovsdb-servers report the status as master to the RA script when the master ip is moved to the node. The issue is fixed upstream.
>
> When the issue occurs, the neutron server will not be able to connect the OVN North and South DB servers and all Create/Update/Delete APIs to the neutron server will fail.
>
> Restarting the ovn-dbs-bundle resource will resolve the issue. Run the below command in one of the controller node:
>
> "pcs resource restart ovn-dbs-bundle"

BZ#1579417

> SNAT support requires configuring VXLAN tunnels regardless of the encapsulation used in the tenant networks. It is also necessary to

> configure the MTU correctly when using VLAN tenant networks, since the VXLAN Tunnel header is added to the payload and this could cause the packet to exceed the default MTU (1500 Bytes).
>
> The VXLAN tunnels have to be properly configured in order for the SNAT traffic to flow through them.
> When using VLAN tenant networks, use one of the following methods to configure MTU so that SNAT traffic can flow through the VXLAN tunnels::
> * Configure VLAN tenant based networks to use an MTU of 1450 on a per network configuration.
> * Set NeutronGlobalPhysnetMtu heat parameter to 1450. Note: the implication of this means all flat/VLAN provider networks will have a 1450 MTU, which may not be desirable (especially for external provider networks).
> * Configure tenant network underlay with MTU of 1550 (or higher). This includes setting the MTU in the NIC templates for tenant network NIC.

BZ#1581337

> In order to use the PING type health monitor, the HAProxy (default software we use in our driver for network load balancing) version must be at least 1.6. Any use of an older HAProxy version makes the health-check be TCP connect without the user's knowledge.
>
> The upstream community fixed that by adding a check in the code, that determine the HAProxy version that is in use and acts accordingly:
> If HAProxy version 1.6 or later, we can use PING.
> Otherwise, we keep using TCP connect (in the absence of any other solution for those haproxy versions, it is better to do so rather than breaking it altogether).
>
> The problem we have in OSP13 GA is that we ship HAProxy as a part of RHEL channels, which uses an old version of HAProxy. Thus, when OSP13 users configure the PING type health monitor, they will get TCP connect instead.

BZ#1583541

> SRIOV based Compute instances have no connectivity to OVS Compute instances if they are on different networks. The workaround is to use an external router that is connected to both VLAN provider networks.

BZ#1584518

> RHOSP does not configure the availability of DifferentHostFilter / SameHostFilter by default in nova, and these settings are necessary to properly complete some tests. As such, several security group tests might randomly fail.
>
> You should skip those tests, or alternatively add those filters to your nova configuration.

BZ#1584762

> If Telemetry is manually enabled on the undercloud, `hardware.*` metrics does not work due to a misconfiguration of the firewall on each of the

```
nodes.

As a workaround, you need to manually set the `snmpd` subnet with the
control plane network by adding an extra template for the undercloud
deployment as follows"

parameter_defaults:
  SnmpdIpSubnet: 192.168.24.0/24
```

BZ#1588186

```
A race condition causes Open vSwitch to not connect to the Opendaylight
openflowplugin. A fix is currently being implemented for a 13.z release of
this product.
```

BZ#1590114

```
If Telemetry is manually enabled on the undercloud, `hardware.*` metrics
does not work due to a misconfiguration of the firewall on each of the
nodes.

As a workaround, you need to manually set the `snmpd` subnet with the
control plane network by adding an extra template for the undercloud
deployment as follows"

parameter_defaults:
  SnmpdIpSubnet: 192.168.24.0/24
```

BZ#1590560

```
The ceph-ansible utility does not always remove the ceph-create-keys
container from the same node where it was created.

Because of this, the deployment may fail with the message "Error response
from daemon: No such container: ceph-create-keys." This may affect any
ceph-ansible run, including fresh deployments, that have:
* multiple compute notes or
* a custom role behaving as ceph client which is also hosting a service
consuming ceph.
```

BZ#1590938

```
If you deploy more than three OSDs on RHCS3 and set the PG number for your
pools as determined by pgcalc (https://access.redhat.com/labs/cephpgc),
deployment will fail because ceph-ansible creates pools before all OSDs
are active.

To avoid the problem, set the default PG number to 32 and when the
deployment is finished, manually raise the PG number as described in the
Storage Strategies Guide, https://access.redhat.com/documentation/en-
us/red_hat_ceph_storage/3/html/storage_strategies_guide/placement_groups_p
gs#set_the_number_of_pgs.
```

BZ#1590939

```
Because ceph-ansible OpenStack pool tasks have an incorrect container
name, it is not yet possible to colocate Ceph MONs and OSDs.
Standard HCI (Computes + OSDs) is not affected.
```

BZ#1593290

```
After restarting the nova-compute service when a guest with SR-IOV-based
network interface(s) attached is running and removing the guest, it is no
longer possible to attach SR-IOV VFs on that node to any guest. This is
because available devices are enumerated on service startup but as the
device is attached to a guest it is not included in the list of host
devices.

You must restart the 'nova-compute' service after removing the guest.
After removing the guest and restarting the service, the list of available
SR-IOV devices will be correct.
```

BZ#1593715

```
Insecure registry list is being updated later than some container images
are pulled during a major upgrade. As such, container images from newly
introduced insecure registry fails to download during `openstack overcloud
upgrade run` command.

You can use one of the following workarounds:

Option A: Update the /etc/sysconfig/docker file manually on nodes which
have containers managed by Pacemaker, and add any newly introduced
insecure registries.

Option B: run `openstack overcloud deploy` command right before upgrading,
and provide the desired new insecure registry list using an environment
file with the DockerInsecureRegistryAddress parameter.

All container images should download successfully during upgrade.
```

BZ#1593757

```
Enabling Octavia on an existing overcloud deployment reports as a success,
but the Octavia API endpoints are not reachable because the firewall rules
on the Controller nodes are misconfigured.

Workaround:

On all controller nodes, add firewall rules and make sure they are inserted
before the DROP rule:

IPv4:
  # iptables -A INPUT -p tcp -m multiport --dports 9876 -m state --state
NEW -m comment --comment "100 octavia_api_haproxy ipv4" -j ACCEPT
  # iptables -A INPUT -p tcp -m multiport --dports 13876 -m state --state
NEW -m comment --comment "100 octavia_api_haproxy_ssl ipv4" -j ACCEPT
  # iptables -A INPUT -p tcp -m multiport --dports 9876,13876 -m state --
state NEW -m comment --comment "120 octavia_api ipv4" -j ACCEPT
```

```
IPv6:
  # ip6tables -A INPUT -p tcp -m multiport --dports 9876 -m state --state
NEW -m comment --comment "100 octavia_api_haproxy ipv6" -j ACCEPT
  # ip6tables -A INPUT -p tcp -m multiport --dports 13876 -m state --state
NEW -m comment --comment "100 octavia_api_haproxy_ssl ipv6" -j ACCEPT
  # ip6tables -A INPUT -p tcp -m multiport --dports 9876,13876 -m state --
state NEW -m comment --comment "120 octavia_api ipv6" -j ACCEPT

Restart HAProxy:
  # docker restart haproxy-bundle-docker-0
```

BZ#1595363

```
During the fast forward upgrade process, users upgrade the undercloud from
version 10 to version 11. In some situations, the nova-api.log might
report the following error:

`Unexpected API Error. Table 'nova_cell0.instances' doesn't exist`

You can resolve this error by running the following command:

$ sudo nova-manage api_db sync

This issue is non-critical and should not impede the fast forward upgrade
process in a major way.
```

3.2. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE 19 JULY 2018

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#1592823

```
Logs from Ansible playbooks now include timestamps that provide
information about the timing of actions during deployment, updates, and
upgrades.
```

3.2.2. 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#1578312

```
When the OVSDB server fails over to a different controller node, a
```

> reconnection from neutron-server/metadata-agent does not take place because they are not detecting this condition.
>
> As a result, booting VMs may not work as metadata-agent will not provision new metadata namespaces and the clustering is not behaving as expected.
>
> A possible workaround is to restart the ovn_metadata_agent container in all the compute nodes after a new controller has been promoted as master for OVN databases. Also increase the ovsdb_probe_interval on the plugin.ini to a value of 600000 milliseconds.

3.2.3. Known Issues

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

BZ#1515815

> When the router gateway is cleared, the Layer 3 flows related to learned IP addresses is not removed. The learned IP addresses include the PNF and external gateway IP addresses. This leads stale flows, but not any functional issue. The external gateway and IP address does not change frequently. The stale flows will be removed when the external network is deleted.

BZ#1519783

> Neutron may issue an error claiming that the Quota has been exceed for Neutron Router creation. This is a known issue where multiple router resources are created with a single create request in Neutron DB due to a bug with networking-odl. The workaround for this issue is to delete the duplicated routers using the OpenStack Neutron CLI and create a router again, resulting with a single instance.

BZ#1559055

> OpenDaylight logging might be missing earlier logs. This is a known issue with journald logging of OpenDaylight (using the “docker logs opendaylght_api” command). The current workaround is to switch OpenDaylight logging to the “file” mechanism which will log inside of the container to /opt/opendaylight/data/logs/karaf.log. To do this, configure the following heat parameter: OpenDaylightLogMechanism: ‘file’.

BZ#1568311

> Layer 3 connectivity between nova instances across multiple subnets may fail when an instance without a floating IP tries to reach another instance that has a floating IP on another router. This occurs when nova instances are spread across multiple compute nodes. There is no suitable workaround for this issue.

BZ#1568976

> During deployment, one or more OpenDaylight instances may fail to start correctly due to a feature loading bug. This may lead to a deployment or

> functional failure.
>
> When a deployment passes, only two of the three OpenDaylight instances must be functional for the deployment to succeed. It is possible that the third OpenDaylight instance started incorrectly. Check the health status of each container with the `docker ps` command. If it is unhealthy, restart the container with `docker restart opendaylight_api`.
>
> When a deployment fails, the only option is to restart the deployment. For TLS-based deployments, all OpenDaylight instances must boot correctly or deployment will fail.

BZ#1583541

> SRIOV based Compute instances have no connectivity to OVS Compute instances if they are on different networks. The workaround is to use an external router that is connected to both VLAN provider networks.

BZ#1588186

> A race condition causes Open vSwitch to not connect to the Opendaylight openflowplugin. A fix is currently being implemented for a 13.z release of this product.

BZ#1593757

> Enabling Octavia on an existing overcloud deployment reports as a success, but the Octavia API endpoints are not reachable because the firewall rules on the Controller nodes are misconfigured.
>
> Workaround:
>
> On all controller nodes, add firewall rules and make sure they are inserted before the DROP rule:
>
> IPv4:
>
> ```
>   # iptables -A INPUT -p tcp -m multiport --dports 9876 -m state --state
> NEW -m comment --comment "100 octavia_api_haproxy ipv4" -j ACCEPT
>   # iptables -A INPUT -p tcp -m multiport --dports 13876 -m state --state
> NEW -m comment --comment "100 octavia_api_haproxy_ssl ipv4" -j ACCEPT
>   # iptables -A INPUT -p tcp -m multiport --dports 9876,13876 -m state --
> state NEW -m comment --comment "120 octavia_api ipv4" -j ACCEPT
> ```
>
> IPv6:
>
> ```
>   # ip6tables -A INPUT -p tcp -m multiport --dports 9876 -m state --state
> NEW -m comment --comment "100 octavia_api_haproxy ipv6" -j ACCEPT
>   # ip6tables -A INPUT -p tcp -m multiport --dports 13876 -m state --state
> NEW -m comment --comment "100 octavia_api_haproxy_ssl ipv6" -j ACCEPT
>   # ip6tables -A INPUT -p tcp -m multiport --dports 9876,13876 -m state --
> state NEW -m comment --comment "120 octavia_api ipv6" -j ACCEPT
> ```
>
> Restart HAProxy:
>
> ```
>   # docker restart haproxy-bundle-docker-0
> ```

3.3. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE 29 AUGUST 2018

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#1561961

```
This feature adds support for PCI device NUMA affinity policies. These are
configured as part of the `[pci]alias` configuration options. There are
three policies supported:
- `required`
- `legacy`
- `preferred`
In all cases, strict NUMA affinity is provided if possible. The key
difference between the policies is how much NUMA affinity you can forsake
before failing to schedule.
These policies allow you to configure how strict your NUMA affinity is on
a per-device basis or, more specifically, per device alias. This is useful
to ensure maximum resource utilization.
When the 'preferred' policy is configured for a PCI device, nova now
utilizes CPUs on a different NUMA node from the NUMA node of the PCI
device if this is all that is available. This results in increased
resource utilization with the downside of reduced performance for these
instances.
```

BZ#1564918

```
Previously, Ironic considered just one IPMI error as retryable. That might
have caused unjustified Ironic failure. With this enhancement, Ironic
treats more types of IPMI error messages as retryable by the IPMI-backed
hardware interfaces, such as power and management hardware interfaces.
Specifically, "Node busy", "Timeout", "Out of space", and "BMC
initialization in progress" IPMI errors cause Ironic to retry the IPMI
command. The result is improved reliability of IPMI based communication
with BMC.
```

BZ#1571741

```
Nova's libvirt driver now allows the specification of granular CPU feature
flags when configuring CPU models.

One benefit of this change is the alleviation of a performance degradation
experienced on guests running with certain Intel-based virtual CPU models
after application of the "Meltdown" CVE fixes.  This guest performance
impact is reduced by exposing the CPU feature flag 'PCID' ("Process-
Context ID") to the *guest* CPU, assuming that the PCID flag is available
in the physical hardware itself.
```

```
For more details, refer to the  documentation of
``[libvirt]/cpu_model_extra_flags`` in ``nova.conf`` for usage details.
```

BZ#1574349

```
It is possible to create the stonith resources for the cluster
automatically before the overcloud deployment.
Before the start of the deployment, run the following command:
openstack overcloud generate fencing --ipmi-lanplus --output
/home/stack/fencing.yaml /home/stack/instackenv.json

Then pass '-e /home/stack/fencing.yaml' to the list of arguments to the
deploy command. This creates the necessary stonith resources for the
cluster automatically.
```

BZ#1578633

```
rhosp-director-images are now multi-arch. OSP 13 now has overcloud full and
ironic python agent images for ppc64le. The resulting rhosp-director-
images were adjusted to accommodate this change.
As a result, rhosp-director-images and rhosp-director-images-ipa are now
meta-packages, with rhosp-director-images-<arch> and rhosp-director-images-
ipa-<arch> rpms added for multi-arch support.
```

BZ#1578636

```
rhosp-director-images are now multi-arch. OSP 13 now has overcloud full and
ironic python agent images for ppc64le. The resulting rhosp-director-
images were adjusted to accommodate this change.
As a result, rhosp-director-images and rhosp-director-images-ipa are now
meta-packages, with rhosp-director-images-<arch> and rhosp-director-images-
ipa-<arch> rpms added for multi-arch support.
```

BZ#1579691

```
Nova's libvirt driver now allows the specification of granular CPU feature
flags when configuring CPU models.
One benefit of this is the alleviation of a performance degradation
experienced on guests running with certain Intel-based virtual CPU models
after application of the "Meltdown" CVE fixes. This guest performance
impact is reduced by exposing the CPU feature flag 'PCID' ("Process-
Context ID") to the *guest* CPU, assuming that the PCID flag is available
in the physical hardware itself.
This change removes the restriction of having only 'PCID' as the only CPU
feature flag and allows for the addition and removal of multiple CPU flags,
making way for other use cases.
For more information, refer to the  documentation of
``[libvirt]/cpu_model_extra_flags`` in ``nova.conf``.
```

BZ#1601472

```
The procedures for upgrading from RHOSP 10 to RHOSP 13 with NFV deployed
have been retested and updated for DPDK and SR-IOV environments.
```

BZ#1606224

> With this update, Ceph storage is supported by KVM virtualization on all CPU architectures supported by Red Hat.

BZ#1609352

> This enhancement sees the addition of GA containers for nova and utilities, and Technology Preview containers for Cinder, Glance, Keystone, Neutron, and Swift on IBM Power LE.

BZ#1619311

> rhosp-director-images are now multi-arch. OSP 13 now has overcloud full and ironic python agent images for ppc64le. The resulting rhosp-director-images were adjusted to accommodate this change.
> As a result, rhosp-director-images and rhosp-director-images-ipa are now meta-packages, with rhosp-director-images-<arch> and rhosp-director-images-ipa-<arch> rpms added for multi-arch support.

3.3.2. 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#1523864

> This update adds support for use of Manila IPv6 export locations and access rules with Dell-EMC Unity and VNX back ends.

BZ#1549770

> Containers are now the default deployment method. There is still a way to deploy the baremetal services in environments/baremetal-services.yaml, but this is expected to eventually disappear.
>
> Environment files with resource registries referencing environments/services-docker must be altered to the environments/services paths. If you need to retain any of the
> deployed baremetal services, update references to environments/services-baremetal instead of the originally placed environments/services.

BZ#1565028

> README has been added to /var/log/opendaylight, stating the correct OpenDaylight log path.

BZ#1570039

> The compress option for the containerized logrotate service to compress rotated logs by default has been added. The delaycompress option ensures the first rotation of a log file remains uncompressed.

BZ#1575752

> In previous versions, the *NetName parameters (e.g. InternalApiNetName) changed the names of the default networks. This is no longer supported. To change the names of the default networks, use a custom composable network file (network_data.yaml) and include it with your 'openstack overcloud deploy' command using the '-n' option. In this file, set the "name_lower" field to the custom net name for the network you want to change. For more information, see "Using Composable Networks" in the Advanced Overcloud Customization guide.
> In addition, you need to add a local parameter for the ServiceNetMap table to network_environment.yaml and override all the default values for the old network name to the new custom name. You can find the default values in /usr/share/openstack-tripleo-heat-templates/network/service_net_map.j2.yaml. This requirement to modify ServiceNetMap will not be necessary in future OSP-13 releases.

BZ#1592528

> In rare circumstances, after rebooting controller nodes several times, RabbitMQ may be running in an inconsistent state that will block API operations on the overcloud.
>
> The symptoms for this issue are:
> - Entries in any of the OpenStack service logs of the form:
> DuplicateMessageError: Found duplicate message(629ff0024219488499b0fac0cacaa3a5). Skipping it.
> - "openstack network agent list" returns that some agents are DOWN
>
> To restore normal operation, run the following command on any of the controller nodes (you only need to do this on one controller):
> pcs resource restart rabbitmq-bundle

3.3.3. Known Issues

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

BZ#1557794

> A regression was identified in the procedure for backing up and restoring the director undercloud. As a result, the procedure requires modification and verification before it can be published.
>
> The book 'Back Up and Restore the Director Undercloud' is therefore not available with the general availability of Red Hat OpenStack Platform 13. The procedure will be updated as a priority after the general availability release, and published as soon as it is verified.

BZ#1579025

> OVN pacemaker Resource Agent (RA) script sometimes does not handle the promotion action properly when pacemaker tries to promote a slave node. This is seen when the ovsdb-servers report the status as master to the RA script when the master ip is moved to the node. The issue is fixed

```
upstream.

When the issue occurs, the neutron server will not be able to connect the
OVN North and South DB servers and all Create/Update/Delete APIs to the
neutron server will fail.

Restarting the ovn-dbs-bundle resource will resolve the issue. Run the
below command in one of the controller node:

"pcs resource restart ovn-dbs-bundle"
```

BZ#1584762

```
If Telemetry is manually enabled on the undercloud, `hardware.*` metrics
does not work due to a misconfiguration of the firewall on each of the
nodes. As a workaround, you need to manually set the `snmpd` subnet with
the control plane network by adding an extra template for the undercloud
deployment as follows:
parameter_defaults:
  SnmpdIpSubnet: 192.168.24.0/24
```

CHAPTER 4. TECHNICAL NOTES

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

4.1. RHEA-2018:2086 — RED HAT OPENSTACK PLATFORM 13.0 ENHANCEMENT ADVISORY

The bugs contained in this section are addressed by advisory RHEA-2018:2086. Further information about this advisory is available at link: https://access.redhat.com/errata/RHEA-2018:2086.

ceph-ansible

BZ#1590560

> The ceph-ansible utility does not always remove the ceph-create-keys container from the same node where it was created.
>
> Because of this, the deployment may fail with the message "Error response from daemon: No such container: ceph-create-keys." This may affect any ceph-ansible run, including fresh deployments, that have:
> * multiple compute notes or
> * a custom role behaving as ceph client which is also hosting a service consuming ceph.

gnocchi

BZ#1533206

> The openstack-gnocchi packages have been renamed to gnocchi. The openstack- prefix was removed because of an upstream project scoping change. Gnocchi has been moved out of the OpenStack umbrella and is maintained as a stand-alone project.

opendaylight

BZ#1568012

> Connecting to an external IP fails when associating a floating IP to an instance then disassociating the floating IP. This situation happens in a tenant VLAN network when:
> * a VM spawned on a non-NAPT switch is associated with a floating IP and
> * the floating IP is removed.
> This results in a missing flow (sporadically) in the FIB table of NAPT switch.
>
> Due to the missing FIB table entry, the VM loses connectivity to the public network.
>
> Associating the floating IP to the VM restores connectivity to the public network. As long as the floating IP is associated with the VM, it will be able to connect to the internet. However, you will lose a public IP/floating IP from the external network.

openstack-cinder

BZ#1557331

> Previously, the cinder service had to be restarted twice when performing an offline upgrade because of the rolling upgrade mechanism.
>
> The double system restart can be skipped with the new optional parameter -called "--bump-versions"- added to the cinder-manage db sync command.

BZ#1572220

> The Block Storage service (cinder) uses a synchronization lock to prevent duplicate entries in the volume image cache. The scope of the lock was too broad and caused simultaneous requests to create a volume from an image to compete for the lock, even when the image cache was not enabled.
>
> These simultaneous requests to create a volume from an image would be serialized and not run in parallel.
>
> As a result, the synchronization lock has been updated to minimize the scope of the lock and to take effect only when the volume image cache is enabled.
>
> Now, simultaneous requests to create a volume from an image run in parallel when the volume image cache is disabled. When the volume image cache is enabled, locking is minimized to ensure only a single entry is created in the cache.

openstack-manila

BZ#1468020

> The Shared File System service (manila) now provides IPv6 access rule support with NetApp ONTAP cDOT driver, which lets you use manila with IPv6 environments.
>
> As a result, the Shared File System service now supports exporting shares backed by NetApp ONTAP back ends over IPv6 networks. Access to the exported shares is controlled by IPv6 client addresses.

BZ#1469208

> The Shared File System service (manila) supports mounting shared file systems backed by a Ceph File System (CephFS) via the NFSv4 protocol. NFS-Ganesha servers operating on Controller nodes are used to export CephFS to tenants with high availability (HA). Tenants are isolated from one another and may only access CephFS through the provided NFS gateway interface. This new feature is fully integrated into director, enabling CephFS back end deployment and configuration for the Shared File System service.

openstack-neutron

BZ#1552108

> When an interface is added or removed to or from a router and isolated metadata is enabled on the DHCP Agent, the metadata proxy for that network is not updated.
>
> As such, instances would not able to fetch metadata if they are on a network which is not connected to a router.
>
> You need to update metadata proxies when a router interface is added or removed. The instances will then be able to fetch metadata from the DHCP namespace when their networks become isolated.

openstack-selinux

BZ#1561711

> Previously, the virtlogd service logged redundant AVC denial errors when a guest virtual machine was started. With this update, the virtlogd service no longer attempts to send shutdown inhibition calls to systemd, which prevents the described errors from occurring.

openstack-swift

BZ#1419556

> The Object Store service (swift) can now integrate with Barbican to transparently encrypt and decrypt your stored (at-rest) objects. At-rest encryption is distinct from in-transit encryption and refers to the objects being encrypted while being stored on disk.
>
> Swift objects are stored as clear text on disk. These disks can pose a security risk if not properly disposed of when they reach end-of-life. Encrypting the objects mitigates that risk.
>
> Swift performs these encryption tasks transparently, with the objects being automatically encrypted when uploaded to swift, then automatically decrypted when served to a user. This encryption and decryption is done using the same (symmetric) key, which is stored in Barbican.

openstack-tripleo-common

BZ#1560422

> Octavia does not scale to practical workloads because the default configured quotas for the "service" project limits the number of Octavia load balancers that can be created in the overcloud.
>
> To mitigate this problem, as the overcloud admin user, set the required quotas to unlimited or some sufficiently large value. For example, run the following commands on the undercloud:
>
> ```
> # source ~/overcloudrc
> #  openstack quota set --cores -1 --ram -1 --ports -1 --instances -1 --
> ```

```
secgroups -1 service
```

BZ#1588838

```
The tripleo.plan_management.v1.update_roles workflow did not pass the
overcloud plan name (swift container name) or zaqar queue name to the sub-
workflow it triggered. This caused incorrect behaviour when using an
overcloud plan name other than the default ('overcloud'). This fix
correctly passes these parameters and restores the correct behaviour.
```

BZ#1566463

```
The 'docker kill' command does not exit if the container is set to
automatically restart. If a user attempts to run 'docker kill
<container>', it may hang indefinitely.  In this case, CTRL+C will stop
the command.

To avoid the problem, use 'docker stop' (instead of 'docker kill') to stop
a containerized service.
```

BZ#1452979

```
Cause: The "openstack overcloud node configure" command would only take
image names not image IDs for "deploy-kernel" and "deploy-ramdisk"
parameters. Image IDs are now accepted after this fix.
```

openstack-tripleo-heat-templates

BZ#1341176

```
This enhancement adds support for deploying RT enabled compute nodes from
director alongside "regular" compute nodes.

1. Based on tripleo-heat-templates/environments/compute-real-time-
example.yaml, create a compute-real-time.yaml environment file that sets
the parameters for the ComputeRealTime role with at least the correct
values for:

 * IsolCpusList and NovaVcpuPinSet: a list of CPU cores that should be
reserved for real-time workloads. This depends on your CPU hardware on
your real-time compute nodes.

 * KernelArgs: set to "default_hugepagesz=1G hugepagesz=1G hugepages=X"
with X depending on the number of guests and how much memory they will
have.

2. Build and upload the overcloud-realtime-compute image:

 * Prepare the repos (for CentOS):
   - sudo yum install -y
https://trunk.rdoproject.org/centos7/current/python2-tripleo-repos-
XXX.el7.centos.noarch.rpm
   - sudo -E tripleo-repos current-tripleo-dev
```

```
   - export DIB_YUM_REPO_CONF="/etc/yum.repos.d/delorean*
/etc/yum.repos.d/quickstart*"

 * openstack overcloud image build --image-name overcloud-realtime-compute
--config-file /usr/share/openstack-tripleo-common/image-yaml/overcloud-
realtime-compute.yaml --config-file /usr/share/openstack-tripleo-
common/image-yaml/overcloud-realtime-compute-centos7.yaml

 * openstack overcloud image upload --update-existing --os-image-name
overcloud-realtime-compute.qcow2

3. Create roles_data.yaml with ComputeRealTime and all other required
roles, for example: openstack overcloud roles generate -o
~/rt_roles_data.yaml Controller ComputeRealTime ... and assign the
ComputeRealTime role to the real-time nodes in one of the usual ways. See
https://docs.openstack.org/tripleo-
docs/latest/install/advanced_deployment/custom_roles.html

4. Deploy the overcloud:
openstack overcloud deploy --templates -r ~/rt_roles_data.yaml -e
./tripleo-heat-templates/environments/host-config-and-reboot.yaml -e
./compute-real-time.yaml [...]
```

BZ#1552583

```
The glance-direct method requires a shared staging area when used in a HA
configuration. Image uploads using the 'glance-direct' method may fail in
an HA environment if a common staging area is not present. Incoming
requests to the controller nodes are distributed across the available
controller nodes. One controller handles the first step and another
controller handles the second request with both controllers writing the
image to different staging areas. The second controller will not have
access to the same staging area used by the controller handling the first
step.

Glance supports multiple image import methods, including the 'glance-
direct' method.  This method uses a three-step approach: creating an image
record, uploading the image to a staging area, and then transferring the
image from the staging area to the storage backend so the image becomes
available. In an HA setup (i.e., with 3 controller nodes), the glance-
direct method requires a common staging area using a shared file system
across the controller nodes.

The list of enabled Glance import methods can now be configured. The
default configuration does not enable the 'glance-direct' method (web-
download is enabled by default). To avoid the issue and reliably import
images to Glance in an HA environment, do not enable the 'glance-direct'
method.
```

BZ#1572238

```
The openvswitch systemd script deletes the /run/openvswitch folder when
stopping it in the host.
The /run/openvswitch path inside the ovn-controller container becomes a
```

stale directory. When the service is started again, it recreates the folder. In order for ovn-controller to access this folder again, the folder has to be remounted or the ovn-controller container restarted.

BZ#1309550

A new CinderRbdExtraPools Heat parameter has been added which specifies a list of Ceph pools for use with RBD backends for Cinder. An extra Cinder RBD backend driver is created for each pool in the list. This is in addition to the standard RBD backend driver associated with the CinderRbdPoolName. The new parameter is optional and defaults to an empty list. All of the pools are associated with a single Ceph cluster.

BZ#1518126

Redis is unable to correctly replicate data across nodes in a HA deployment with TLS enabled. Redis follower nodes will not contain any data from the leader node. It is recommended to disable TLS for Redis deployments.

BZ#1540239

This enhancement adds support for sending metrics data to a Gnocchi DB instance.

The following new parameters for collectd composable service were added. If CollectdGnocchiAuthMode is set to 'simple', then CollectdGnocchiProtocol, CollectdGnocchiServer, CollectdGnocchiPort and CollectdGnocchiUser are taken into account for configuration.

If CollectdGnocchiAuthMode is set to 'keystone', then CollectdGnocchiKeystone* parameters are taken into account for configuration.

Following is a detailed description of added parameters:

```
  CollectdGnocchiAuthMode:
    type: string
    description: >
      Type of authentication Gnocchi server is using. Supported values are
      'simple' and 'keystone'.
    default: 'simple'
  CollectdGnocchiProtocol:
    type: string
    description: API protocol Gnocchi server is using.
    default: 'http'
  CollectdGnocchiServer:
    type: string
    description: >
      The name or address of a gnocchi endpoint to which we should
      send metrics.
    default: nil
  CollectdGnocchiPort:
    type: number
```

```
    description: The port to which we will connect on the Gnocchi server.
    default: 8041
  CollectdGnocchiUser:
    type: string
    description: >
      Username for authenticating to the remote Gnocchi server using
simple
      authentication.
    default: nil
  CollectdGnocchiKeystoneAuthUrl:
    type: string
    description: Keystone endpoint URL to authenticate to.
    default: nil
  CollectdGnocchiKeystoneUserName:
    type: string
    description: Username for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneUserId:
    type: string
    description: User ID for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystonePassword:
    type: string
    description: Password for authenticating to Keystone
    default: nil
  CollectdGnocchiKeystoneProjectId:
    type: string
    description: Project ID for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneProjectName:
    type: string
    description: Project name for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneUserDomainId:
    type: string
    description: User domain ID for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneUserDomainName:
    type: string
    description: User domain name for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneProjectDomainId:
    type: string
    description: Project domain ID for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneProjectDomainName:
    type: string
    description: Project domain name for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneRegionName:
    type: string
    description: Region name for authenticating to Keystone.
    default: nil
  CollectdGnocchiKeystoneInterface:
    type: string
    description: Type of Keystone endpoint to authenticate to.
```

```
    default: nil
  CollectdGnocchiKeystoneEndpoint:
    type: string
    description: >
      Explicitly state Gnocchi server URL if you want to override
      Keystone value
    default: nil
  CollectdGnocchiResourceType:
    type: string
    description: >
      Default resource type created by the collectd-gnocchi plugin in
Gnocchi
      to store hosts.
    default: 'collectd'
  CollectdGnocchiBatchSize:
    type: number
    description: Minimum number of values Gnocchi should batch.
    default: 10
```

BZ#1566376

```
The OVN metadata service was not being deployed in DVR based environment.
Therefore, instances were not able to fetch metadata such as instance
name, public keys, etc.

This patch enables the aforementioned service so that any booted instance
can fetch metadata.
```

BZ#1568120

```
The Heat templates for Cinder backend services were triggering Puppet to
deploy the cinder-volume service on the overcloud host, regardless of
whether the service is meant to be deployed in a container. This caused
the cinder-volume service to be deployed twice: in a container as well as
on the host.

Because of this, the OpenStack volume operations (such as creating and
attaching a volume) would occasionally fail when the operation was handled
by the rogue cinder-volume service running on the host.

As a result, the Cinder backend heat templates have been updated to not
deploy a second instance of the cinder-volume service.
```

BZ#1573597

```
A poorly performing Swift cluster used as a Gnocchi back end can generate
503 errors in the collectd log and "ConnectionError: ('Connection
aborted.', CannotSendRequest())" errors in in gnocchi-metricd.conf.
To mitigate the problem, increase the value of the
CollectdDefaultPollingInterval parameter or improve the Swift cluster
performance.
```

BZ#1575023

> The manila-share service fails to initialize because changes to ceph-ansible's complex ceph-keys processing generate incorrect content in the /etc/ceph/ceph.client.manila.keyring file.
>
> To allow the manila-share service to initialize:
> 1) Make a copy of /usr/share/openstack/tripleo-heat-templates to use for the overcloud deploy.
>
> 2) Edit the .../tripleo-heat-templates/docker/services/ceph-ansible/ceph-base.yaml file to change all triple backslashes in line 295 to single backslashes.
> Before:
> mon_cap: 'allow r, allow command \\\"auth del\\\", allow command \\\"auth caps\\\", allow command \\\"auth get\\\", allow command \\\"auth get-or-create\\\"'
> After:
> mon_cap: 'allow r, allow command \"auth del\", allow command \"auth caps\", allow command \"auth get\", allow command \"auth get-or-create\"'
>
> 3) Deploy the overcloud substituting the path to the copy of tripleo-heat-templates wherever /usr/share/openstack-tripleo-heat templates occurred in your original overcloud-deploy command.
>
> The ceph key /etc/ceph/ceph.client.manila.keyring file will have proper contents and the manila-share service will initialize properly.

BZ#1552214

> When configuring the cinder-volume service for HA, cinder's DEFAULT/host configuration was set to "hostgroup". Other cinder services (cinder-api, cinder-scheduler, cinder-backup) would use "hostgroup" for their configuration, regardless of which overcloud node was running the service. Log messages from these services looked like they all originated from the same "hostgroup" host, which made it difficult to know which node generated the message.
>
> When deploying for HA, cinder-volume's backend_host is set to "hostgroup" instead of setting DEFAULT/host to that value. This ensures each node's DEFAULT/host value is unique.
>
> Consequently, log messages from cinder-api, cinder-scheduler, and cinder-backup are correctly associated with the node that generated the message.

BZ#1578901

> After upgrading to a new release, Block Storage services (cinder) were stuck using the old RPC versions from the prior release. Because of this, all cinder API requests requiring the latest RPC versions failed.
>
> When upgrading to a new release, all cinder RPC versions are updated to match the latest release.

python-cryptography

BZ#1556933

> Since version 2.1, python-cryptography checks that the CNS Names used in certificates are compliant with IDN standards. If the found names do not follow this specification, cryptography will fail to validate the certificate and different errors may be found when using OpenStack command line interface or in OpenStack service logs.

BZ#1571358

> After installing python-cryptography build, the initial import from RDO failed because it was missing Obsoletes. The RHEL 7 build of this package is correct and has right Obsoletes entries.
>
> This fix adds the Obsoletes for python-cryptography.

python-ironic-tests-tempest

BZ#1577982

> A tempest plugin (-tests) rpm installed before the upgrade fails after the OSP Release 13 upgrade. The initial upgrade packaging did not include the epoch commands needed to obsolete the old rpm. The sub-rpm is not shipped in OSP 13, and the Obsoletes in the new plugin rpm didn't correctly Obsolete the right rpm.
>
> To fix the issue, correct the obsoletes or manually uninstall the old -rpm and manually install the replacement plugin python2-*-tests-tempest.

python-networking-ovn

BZ#1433533

> To help maintain consistency between the neutron and OVN databases, configuration changes are internally compared and verified in the backend. Each configuration change is assigned a revision number, and a scheduled task validates all create, update, and delete operations made to the databases.

BZ#1503521

> This version introduces support for internal DNS resolution in networking-ovn. Although there are two know limitations,
> one is bz#1581332 which prevents proper resolution of internal fqdn requests via internal dns.
>
> Please note that the extension is not configured by default by tripleo on the GA release. See bz#1577592 for a workaround.

BZ#1550039

> When a subnet is created without a gateway, no DHCP options were added and

> instances on such subnets are not able to obtain DHCP.
>
> The Metadata/DHCP port is used instead for this purpose so that instances can obtain an IP address. You must enable the metadata service. Instances on subnets without a external gateway are now able to obtain their IP addresses through DHCP via the OVN metadata/DHCP port.

BZ#1562731

> The current L3 HA scheduler was not taking the priorities of the nodes into consideration. Therefore, all gateways were being hosted by the same node and the load was not distributed across candidates.
>
> This fix implements an algorithm to select the least loaded node when scheduling a gateway router. Gateway ports are now being scheduled on the least loaded network node distributing the load evenly across them.

BZ#1563678

> When a subport was reassigned to a different trunk on another hypervisor, it did not get its binding info updated and the subport did not transition to ACTIVE.
>
> This fix clears up the binding info when the subport is removed from the trunk. The subport now transitions to ACTIVE when it is reassigned to another trunk port that resides on a different hypervisor.

python-os-brick

BZ#1550974

> When using iSCSI discovery, the node startup configuration was reset from "automatic" to "default", which caused the services to not be started on reboot. This issue is fixing by restoring all startup values after each discovery.

python-zaqar-tests-tempest

BZ#1546285

> Upgrades were having dependencies issues because the collection of tempest plugins were extracted from openstack-*-tests rpm subpackages during the Queens cycle. However, not all of the packaging had the right combination of Provides and Obsoletes. OSP 13 does not have the -tests (unittest sub-rpms).
>
> When attempting to do upgrades with -tests installed from prior release cause failures due to dependencies issues.
>
> To correct this issue, the Obsoletes for the older version of the -tests rpms they were extracted from have been added back.

4.2. RHSA-2018:2214 — IMPORTANT: OPENSTACK-TRIPLEO-HEAT-TEMPLATES SECURITY UPDATE

The bugs contained in this section are addressed by advisory RHSA-2018:2214. Further information about this advisory is available at link: https://access.redhat.com/errata/RHSA-2018:2214.html.

openstack-tripleo-common

BZ#1592823

> Logs from Ansible playbooks now include timestamps that provide information about the timing of actions during deployment, updates, and upgrades.

openstack-tripleo-heat-templates

BZ#1586171

> Previously, overcloud updates failed due to stale cache in OpenDaylight. With this update, OpenDaylight is stopped and the stale cache is removed before upgrading to a new version. Level 1 updates work with OpenDaylight deployments. Level 2 updates are currently unsupported.

BZ#1593757

> Enabling Octavia on an existing overcloud deployment reports as a success, but the Octavia API endpoints are not reachable because the firewall rules on the Controller nodes are misconfigured.
>
> Workaround:
>
> On all controller nodes, add firewall rules and make sure they are inserted before the DROP rule:
>
> ```
> IPv4:
>   # iptables -A INPUT -p tcp -m multiport --dports 9876 -m state --state
> NEW -m comment --comment "100 octavia_api_haproxy ipv4" -j ACCEPT
>   # iptables -A INPUT -p tcp -m multiport --dports 13876 -m state --state
> NEW -m comment --comment "100 octavia_api_haproxy_ssl ipv4" -j ACCEPT
>   # iptables -A INPUT -p tcp -m multiport --dports 9876,13876 -m state --
> state NEW -m comment --comment "120 octavia_api ipv4" -j ACCEPT
>
>
> IPv6:
>   # ip6tables -A INPUT -p tcp -m multiport --dports 9876 -m state --state
> NEW -m comment --comment "100 octavia_api_haproxy ipv6" -j ACCEPT
>   # ip6tables -A INPUT -p tcp -m multiport --dports 13876 -m state --state
> NEW -m comment --comment "100 octavia_api_haproxy_ssl ipv6" -j ACCEPT
>   # ip6tables -A INPUT -p tcp -m multiport --dports 9876,13876 -m state --
> state NEW -m comment --comment "120 octavia_api ipv6" -j ACCEPT
>
> Restart HAProxy:
>   # docker restart haproxy-bundle-docker-0
> ```

BZ#1559055

> OpenDaylight logging might be missing earlier logs. This is a known issue with journald logging of OpenDaylight (using the “docker logs opendaylght_api” command). The current workaround is to switch OpenDaylight logging to the “file” mechanism which will log inside of the container to /opt/opendaylight/data/logs/karaf.log. To do this, configure the following heat parameter: OpenDaylightLogMechanism: ‘file’.

BZ#1559105

> Rerunning an overcloud deploy command against an existing overcloud failed to trigger a restart of any pacemaker managed resource. For example, when adding a new service to haproxy, haproxy would not restart, rendering the newly configured service unavailable until a manual restart of the haproxy pacemaker resource.
>
> With this update, a configuration change of any pacemaker resource is detected, and the pacemaker resource automatically restarts. Any changes in the configuration of pacemaker managed resources is then reflected in the overcloud.

BZ#1589346

> Service deployment tasks within the minor-update workflow were run twice caused by superfluous entries in the list of playbooks. This update removes the superfluous playbook entries and includes host preparation tasks directly in the updated playbook. Actions in minor version updates run once in the desired order.

BZ#1592424

> Previously, the UpgradeInitCommonCommand parameter was not present in heat templates used to deploy the overcloud on pre-provisioned servers. The ‘openstack overcloud upgrade prepare’ command would not perform all of the necessary operations, which caused issues during upgrades in some environments.
>
> This update adds UpgradeInitCommonCommand to the templates used for pre-provisioned servers, allowing the ‘openstack overcloud upgrade prepare’ command to perform the necessary actions.

BZ#1594328

> To enhance security, the default OpenDaylightPassword “admin” is now replaced by a randomly generated 16-digit number. You can overwrite the randomly generated password by specifying a password in a heat template:
>
> ```
> $ cat odl_password.yaml
> parameter_defaults:
>   OpenDaylightPassword: admin
> ```
>
> And then pass the file to the overcloud deploy command:

```
openstack overcloud deploy <other env files> -e odl_password.yaml
```

puppet-opendaylight

BZ#1594333

```
Previously, the Karaf shell (the management shell for OpenDaylight) was
not bound to a specific IP on port 8101, causing the Karaf shell to listen
on the public-facing, external network. This created a security
vulnerability, because the external network could be used to access
OpenDaylight on the port.

This update binds the Karaf shell to the internal API network IP during
deployment, which makes the Karaf shell only accessible on the private
internal API network.
```

4.3. RHBA-2018:2215 — OPENSTACK-NEUTRON BUG FIX ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2018:2215. Further information about this advisory is available at link: https://access.redhat.com/errata/RHBA-2018:2215.html.

opendaylight

BZ#1568311

```
Layer 3 connectivity between nova instances across multiple subnets may
fail when an instance without a floating IP tries to reach another
instance that has a floating IP on another router. This occurs when nova
instances are spread across multiple compute nodes. There is no suitable
workaround for this issue.
```

BZ#1568976

```
During deployment, one or more OpenDaylight instances may fail to start
correctly due to a feature loading bug. This may lead to a deployment or
functional failure.

When a deployment passes, only two of the three OpenDaylight instances
must be functional for the deployment to succeed. It is possible that the
third OpenDaylight instance started incorrectly. Check the health status
of each container with the docker ps command. If it is unhealthy, restart
the container with docker restart opendaylight_api.

When a deployment fails, the only option is to restart the deployment. For
TLS-based deployments, all OpenDaylight instances must boot correctly or
deployment will fail.
```

BZ#1586169

```
Missing parameters from createFibEntry generate a Null Pointer Exception
(NPE) during NAT setup. This bug may result in missing FIB entries from the
routing table, causing NAT or routing to fail. This update adds the proper
```

> parameters to the RPC call. NPE is no longer seen in the OpenDaylight log, and NAT and routing function correctly.

BZ#1587967

> When the NAPT switch is selected on a node without any port in a VLAN network, all flows required are not programmed. External connectivity fails for all VMs in the network that don’t have floating IP addresses. This update adds a pseudo port to create a VLAN footprint in the NAPT switch for VLANs that are part of the router. External connectivity works for VMs without floating IP addresses.

BZ#1588186

> A race condition causes Open vSwitch to not connect to the Opendaylight openflowplugin. A fix is currently being implemented for a 13.z release of this product.

BZ#1515815

> When the router gateway is cleared, the Layer 3 flows related to learned IP addresses is not removed. The learned IP addresses include the PNF and external gateway IP addresses. This leads stale flows, but not any functional issue. The external gateway and IP address does not change frequently. The stale flows will be removed when the external network is deleted.

openstack-neutron

BZ#1591206

> A new configuration option called bridge_mac_table_size has been added for the neutron OVS agent. This value is set as the "other_config:mac-table-size" option on each bridge managed by the openvswitch-neutron-agent. The value controls the maximum number of MAC addresses that can be learned on a bridge. The default value for this new option is 50,000, which should be enough for most systems. Values outside a reasonable range (10 to 1,000,000) will be forced by OVS.

python-networking-odl

BZ#1519783

> Neutron may issue an error claiming that the Quota has been exceed for Neutron Router creation. This is a known issue where multiple router resources are created with a single create request in Neutron DB due to a bug with networking-odl. The workaround for this issue is to delete the duplicated routers using the OpenStack Neutron CLI and create a router again, resulting with a single instance.

python-networking-ovn

BZ#1578312

```
When the OVSDB server fails over to a different controller node, a
reconnection from neutron-server/metadata-agent does not take place
because they are not detecting this condition.

As a result, booting VMs may not work as metadata-agent will not provision
new metadata namespaces and the clustering is not behaving as expected.

A possible workaround is to restart the ovn_metadata_agent container in
all the compute nodes after a new controller has been promoted as master
for OVN databases. Also increase  the ovsdb_probe_interval on the
plugin.ini to a value of 600000 milliseconds.
```

BZ#1582512

```
If the 'dns_nameservers' field is not set for a subnet, the VMs attached to
the subnet have empty /etc/resolv.conf. With this fix, neutron-server gets
the DNS resolver from the /etc/resolv.conf of the host from which it runs
and uses it as the default dns_nameservers for the tenant VMs.
```

4.4. RHBA-2018:2573 — OPENSTACK PLATFORM 13 BUG FIX AND ENHANCEMENT ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2018:2573. Further information about this advisory is available at link: https://access.redhat.com/errata/RHBA-2018:2573

openstack-kuryr-kubernetes

BZ#1585237

```
The controller does not support Nodeport services, and users should not
create them. Nonetheless, Nodeport services are present in some
configurations, and their presence has caused the controller to crash.
To safeguard against such crashes, the controller now ignores Nodeport
services.
```

openstack-manila

BZ#1523864

```
This update adds support for use of Manila IPv6 export locations and access
rules with Dell-EMC Unity and VNX back ends.
```

openstack-manila-ui

BZ#1554935

```
Configuration files for manila-ui plugin were not being copied.  As a
result, the manila panel did not show up on the dashboard.
The instructions for copying all of the configuration files for manila-ui
```

> to the required locations are now present. The manila panel is visible when the user enables the dashboard.

openvswitch

BZ#1551016

> The creation time of OVN ports grew linearly as ports were created. The creation time now remains constant, regardless of the number of ports in the cloud.

python-eventlet

BZ#1607967

> There was an issue in python-eventlet UDP address handling that resulted in some IPv6 addresses being handled incorrectly in some cases. As a result, when receiving DNS responses via UDP, python-eventlet ignored the response and stalled for several seconds, severely impacting performance. This issue is now resolved.

BZ#1612971

> Due to a bug in eventlet, systems that did not configure any nameservers (or in which the nameservrs were unreachable) and that relied only on hosts file for name resolution hit a delay when booting instances. This is because of an attempt to resolve the IPv6 entry even when only an IPv4 host was specified.
> With this fix, eventlet returns immediately without attempting to use network resolution if at least one of the entries is present in the hosts file.

python-oslo-policy

BZ#1600137

> Previously, every time a policy check was made in neutron, the policy file was reloaded and re-evaluated. The re-evaluation of the policy file slowed down API operations substantially for non-admin users.
> With this update, the state of the policy file is saved so the file only reloads if the rules have changed. Neutron API operations for non-admin users are resolved quickly.

python-proliantutils

BZ#1578581

> Because of issues with multiple Sushy object creation on HP Gen10 servers, HPE Gen10 servers were not providing consistent response when accessing the system with id /redfish/v1/Systems/1.
> Instead of using session-based authentication, which is the default authentication method in Sushy, use basic authentication at the time of Sushy object creation. This resolves power request issues.

BZ#1580480

```
When the ironic-dbsync utility tried to load the ironic drivers and when a
driver imported the proliantutils.ilo client module, the proliantutils
library tried to load all of the pysnmp MIBs. If the ironic-dbsync process
resided in an unreadable CWD, pysnmp failed when trying to search for MIBs
in CWD. This resulted in the following error messages in ironic-dbsync.log
on deployment: Unable to load classic driver fake_drac: MIB file
pysnmp_mibs/CPQIDA-MIB.pyc access error: [Errno 13] Permission denied:
'pysnmp_mibs': MibLoadError: MIB file pysnmp_mibs/CPQIDA-MIB.pyc access
error: [Errno 13] Permission denied: 'pysnmp_mibs'
An update to proliantutils ensures that pysnmp does not load all MIBs on
module import. This avoids the situation when an MIB search is attempted
prior to the moment of being explicitly requested by the application.
```

rhosp-release

BZ#1563435

```
When removing older image packages, the post scriptlets sometimes
incorrectly updated the symlinks for image packages.
The scriptlets have been updated to call a script that can be used to fix
the symlinks.
```

4.5. RHBA-2018:2574 — OPENSTACK DIRECTOR BUG FIX ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2018:2574. Further information about this advisory is available at link: https://access.redhat.com/errata/RHBA-2018:2574

instack-undercloud

BZ#1572257

```
Red Hat OpenStack undercloud upgrade failed when the overcloud was in a
Failed state. It failed very late with a cryptic error when trying to
migrate the overcloud stack to use convergence architecture in the post-
configuration step of the upgrade process.
Now, it fails fast and does not allow undercloud upgrade to proceed. The
user receives an error at the beginning of undercloud upgrade. The user
must ensure that the overcloud is in *_COMPLETE state before proceeding
with the undercloud upgrade.
```

BZ#1584666

```
Previously, when the parameter local_mtu was set to 1900 and was specified
in undercloud.conf, the undercloud installation failed.  If the value of
local_mtu was greater than 1500, the undercloud installation failed.
Set global_physnet_mtu to local_mtu. Undercloud installation succeeds when
the value of local_mtu is greater than 1500.
```

BZ#1608173

> Sometimes an undercloud that has SSL enabled failed during installation with the following error: ERROR: epmd error. Failure occurred because the VIP matching the hostname was configured by keepalived after rabbitmq. Ensure that you configure keepalived before rabbitmq. This prevents undercloud installation failure.

openstack-tripleo

BZ#1601472

> The procedures for upgrading from RHOSP 10 to RHOSP 13 with NFV deployed have been retested and updated for DPDK and SR-IOV environments.

openstack-tripleo-common

BZ#1594279

> The 'openstack undercloud backup' command did not capture extended attributes. This caused metadata loss from the undercloud Swift storage object, rendering them unusable. This fix adds the '--xattrs' flag when creating the backup archive. Undercloud Swift storage objects now retain their extended attributes during backup.

BZ#1596763

> When the undercloud imported bare metal nodes from the instackenv.json file and while the UCS driver was being configured, ironic nodes that only differ in pm_service_profile (or ucs_service_profile) fields overrode one another in ironic configuration. This resulted in just one of such ironic nodes ending up in the ironic configuration.
> An update to openstack-tripleo-common ensures that ironic nodes that only differ in pm_service_profile (or ucs_service_profile) fields are still considered distinct.
> All of the ironic nodes that only differ in pm_service_profile or ucs_service_profile fields get imported into ironic.

BZ#1574349

> It is possible to create the stonith resources for the cluster automatically before the overcloud deployment.
> Before the start of the deployment, run the following command:
> openstack overcloud generate fencing --ipmi-lanplus --output /home/stack/fencing.yaml /home/stack/instackenv.json
>
> Then pass '-e /home/stack/fencing.yaml' to the list of arguments to the deploy command. This creates the necessary stonith resources for the cluster automatically.

BZ#1575623

> The Derived Parameters workflow now supports the use of SchedulerHints to identify overcloud nodes.

> Previously, the workflow could not use use SchedulerHints to identify overcloud nodes associated with the corresponding TripleO overcloud role. This caused the overcloud deployment to fail. SchedulerHints support prevents these failures.

BZ#1577853

> The docker healthcheck for OpenDaylight ensured only that the REST interface and neutron NB component was healthy in OpenDaylight. The healthcheck did not include all loaded OpenDaylight components and therefore was not accurate.
> Use diagstatus URI with docker healthcheck to check all of the loaded OpenDaylight components.
> OpenDaylight docker container health status is now more accurate.

openstack-tripleo-heat-templates

BZ#1597379

> The manila-share service container failed to bind-mount PKI trust stores from the controller host. As a result, connections from the manila-share service to the storage back end could not be encrypted using SSL.
> Bind-mount the PKI trust stores from the controller host into the manila-share service container. The connections from the manila-share service to the storage back end can now be encrypted using SSL.

BZ#1597541

> A change in the libvirtd live-migration port range prevents live-migration failures. Previously, libvirtd live-migration used ports 49152 to 49215, as specified in the qemu.conf file. On Linux, this range is a subset of the ephemeral port range 32768 to 61000. Any port in the ephemeral range can be consumed by any other service as well.
> As a result, live-migration failed with the error:
> Live Migration failure: internal error: Unable to find an unused port in range 'migration' (49152-49215).
> The new libvirtd live-migration range of 61152-61215 is not in the ephemeral range. The related failures no longer occur.

BZ#1500594

> Previously, when removing the ceph-osd package from the overcloud nodes, the corresponding Ceph product key was not removed. Therefore, the subscription-manager incorrectly reported that the ceph-osd package was still installed. The script that handles the removal of the ceph-osd package now also removes the corresponding Ceph product key. The script that removes the ceph-osd package and product key executes only during the overcloud update procedure. As a result, subscription-manager list no longer reports that the Ceph OSD is installed.

BZ#1549770

> Containers are now the default deployment method. There is still a way to

deploy the baremetal services in environments/baremetal-services.yaml, but this is expected to eventually disappear.

Environment files with resource registries referencing environments/services-docker must be altered to the environments/services paths. If you need to retain any of the deployed baremetal services, update references to environments/services-baremetal instead of the originally placed environments/services.

BZ#1598469

Previously, the code that supports the Fast Forward Upgrade path for Sahara was missing. As a result, not all of the required changes were applied to Sahara services after a Fast Forward Upgrade from 10 to 13. With this update, the issue has been resolved and Sahara services work correctly after a Fast Forward Upgrade.

BZ#1565028

README has been added to /var/log/opendaylight, stating the correct OpenDaylight log path.

BZ#1567511

In CephFS-NFS driver deployments, the NFS-Ganesha server, backed by CephFS, performs dentry, inode, and attribute caching that is also performed by the libcephfs clients.
The NFS-Ganesha server's redundant caching led to a large memory footprint. It also affected cache coherency.
Turn off NFS-Ganesha server's inode, dentry, and attribute caching. This reduces the memory footprint of the NFS-Ganesha server. Cache coherency issues are less probable.

BZ#1567893

TripleO's capabilities-map.yaml referenced Cinder's Netapp backend in an incorrect file location. The UI uses the capabilities map and was unable to access Cinder's Netapp configuration file.
The capabilities-map.yaml has been updated to specify the correct location for Cinder's Netapp configuration. The UI's properties tab for the Cinder Netapp backend functions correctly.

BZ#1574787

Manila configuration manifests for Dell-EMC storage systems (VNX, Unity, and VMAX) had incorrect configuration options. As a result, the overcloud deployment of manila-share service with Dell Storage systems failed.
The Manila configuration manifests for Dell-EMC storage systems (VNX, Unity, and VMAX) have now been fixed. The overcloud deployment of manila-share service with Dell storage systems completes successfully.

BZ#1584762

If Telemetry is manually enabled on the undercloud, hardware.* metrics does not work due to a misconfiguration of the firewall on each of the nodes. As a workaround, you need to manually set the snmpd subnet with the control plane network by adding an extra template for the undercloud deployment as follows:
parameter_defaults:
SnmpdIpSubnet: 192.168.24.0/24

BZ#1589661

On rare occasions, a deployment failed with the following error log from a container:
standard_init_linux.go:178: exec user process caused "text file busy".
To avoid the race and to avoid deployment failure, do not attempt to write out the docker-puppet.sh file multiple times concurrently.

BZ#1590602

When setting the parameter KernelDisableIPv6 to true in order to disable ipv6, the deployment failed with rabbitmq errors because the Erlang Port Mapper Daemon requires that at least the loopback interface support IPv6 in order to initialize correctly.
To ensure successful deployment when disabling ipv6, do not disable IPv6 on the loopback interface.

BZ#1597665

Docker used journald backend rolls over logs based on size. This resulted in the deletion of some of the older OpenDaylight logs. This issue has been resolved by moving to logging to file instead of console where log file size and rollover can be managed by OpenDaylight. As a result, older logs are persistent for a longer duration than before.

BZ#1542493

If you use a non-standard port for RabbitMQ instance that is for monitoring purposes, the sensu-client container reported an unhealthy state due to not reflecting the port value in the container health check.
The port value now shows in the container health check.

BZ#1564519

The default age for purging deleted database records has been corrected so that deleted records are purged from Cinder's database.
Previously, the CinderCronDbPurgeAge value for Cinder's purge cron job used the wrong value and deleted records were not purged from Cinder's DB when they reached the required default age.

BZ#1569515

The single-nic-vlans network templates in TripleO Heat Templates in OSP 13

> contained an incorrect bridge name for Ceph nodes. If the single-nic-vlans templates were used in a previous deployment, upgrades to OSP 13 failed on the Ceph nodes.
> The bridge name br-storage is now used on Ceph nodes in the single-nic-vlans templates, which matches the bridge name from previous versions. Upgrades to OSP 13 on environments using the single-nic-vlans templates are now successful on Ceph nodes.

BZ#1575752

> In previous versions, the *NetName parameters (e.g. InternalApiNetName) changed the names of the default networks. This is no longer supported. To change the names of the default networks, use a custom composable network file (network_data.yaml) and include it with your 'openstack overcloud deploy' command using the '-n' option. In this file, set the "name_lower" field to the custom net name for the network you want to change. For more information, see "Using Composable Networks" in the Advanced Overcloud Customization guide.
> In addition, you need to add a local parameter for the ServiceNetMap table to network_environment.yaml and override all the default values for the old network name to the new custom name. You can find the default values in /usr/share/openstack-tripleo-heat-templates/network/service_net_map.j2.yaml. This requirement to modify ServiceNetMap will not be necessary in future OSP-13 releases.

BZ#1576627

> yaml-nic-config-2-script.py required interactive user input. The script could not be called in a non-interactive manner for automation purposes. A --yes option has been added. yaml-nic-config-2-script.py can now be called with --yes option and the user is not asked for interactive input.

BZ#1593882

> Previously, some versions of the tripleo-heat-templates contained an error in a setting for the Redis VIP port in the environment file fixed-ips-v6.yaml. If the file fixed-ips-v6.yaml was included on the deployment command line after network-isolation-v6.yaml, the Redis service was placed on the Control Plane network rather than the correct IPv6 network.
> With this update, the file environments/fixed-ips-v6.yaml contains the correct reference to network/ports/vip_v6.yaml, instead of network/ports/vip.yaml. The fixed-ips-v6.yaml environment file contains the correct resource registry entries and the Redis VIP will be created with an IPv6 address, regardless of the order of the included environment files.

BZ#1594910

> TripleO's BlockStorage role was not updated when Cinder services migrated from running on the host to running in containers. The cinder-volume service deployed on the BlockStorage host.
> The BlockStorage role has been updated to deploy the cinder-volume service in a container. The cinder-volume service runs correctly in a container.

BZ#1599329

> An overcloud update with Manila configuration changes failed to deploy those changes to the containerized Manila share-service.
> With this fix, the deployment of the changes is now successful.

BZ#1603538

> With shared storage for /var/lib/nova/instances, like nfs, restarting nova_compute on any compute resulted in owner/group change of the instances virtual ephemeral disks and console.log. As a result, instances lost access to their virtual ephemeral disks and stopped working.
> The scripts to modify the ownership of the instance files in /var/lib/nova/instances have been improved.
> There is now no loss in access to the instance files during restart of nova compute.

BZ#1612342

> The TripleO environment files used for deploying Cinder's Netapp backend were out of date and contained incorrect data. This resulted in failed overcloud deployment.
> The Cinder Netapp environment files have been updated and are now correct.
> You can now deploy an overcloud with a Cinder Netapp backend.

BZ#1573787

> Previously, libvirtd live-migration used ports 49152 to 49215, as specified in the qemu.conf file. On Linux, this range is a subset of the ephemeral port range 32768 to 61000. Any port in the ephemeral range can be consumed by any other service as well.
> As a result, live-migration failed with the error:
> Live Migration failure: internal error: Unable to find an unused port in range 'migration' (49152-49215).
> The new libvirtd live-migration range of 61152 to 61215 is not in the ephemeral range.

BZ#1576572

> Previously, if a nic config template contained a blank line followed by a line starting with a comma, the yaml-nic-config-2-script.py did not reset the starting column of the next row. The nic config template converted by the script was invalid and caused a deployment failure.
> With this update, the script correctly sets the value for the column when the blank line is detected. Scripts that have a blank line followed by a line with a comma are converted correctly.

puppet-nova

BZ#1579691

> Nova's libvirt driver now allows the specification of granular CPU feature

> flags when configuring CPU models.
> One benefit of this is the alleviation of a performance degradation experienced on guests running with certain Intel-based virtual CPU models after application of the "Meltdown" CVE fixes. This guest performance impact is reduced by exposing the CPU feature flag 'PCID' ("Process-Context ID") to the guest CPU, assuming that the PCID flag is available in the physical hardware itself.
> This change removes the restriction of having only 'PCID' as the only CPU feature flag and allows for the addition and removal of multiple CPU flags, making way for other use cases.
> For more information, refer to the documentation of [libvirt]/cpu_model_extra_flags in nova.conf.

puppet-opendaylight

BZ#1599805

> OpenDaylight polls OpenFlow (OF) statistics periodically. These statistics are not being used anywhere currently. This affects OpenDaylight performance. You can disable polling of OF statistics to increase OpenDaylight performance.

puppet-tripleo

BZ#1598038

> Instance HA deployments failed due to a race condition, generating an error: Error: unable to get cib.
> The race was a result of pacemaker properties being set on the compute nodes before the pacemaker cluster was fully up and hence failing with the 'unable to get cib' error.
> This fix results in no errors in the deployment when using IHA.

BZ#1564654

> Previously, if you used uppercase letters in the stack name, the deployment failed.
> This update ensures that a stack name with uppercase letters leads to a successful deployment. Specifically, the bootstrap_host scripts inside the containers now convert strings to lowercase and the same happens for pacemaker properties.

BZ#1570039

> The compress option for the containerized logrotate service to compress rotated logs by default has been added. The delaycompress option ensures the first rotation of a log file remains uncompressed.

BZ#1601497

> Previously, configuring empty string values for some deprecated parameters for Cinder's Netapp backend resulted in an invalid configuration for the Cinder driver, causing Cinder's Netapp backend driver to fail during

initialization.
As of this update, empty string values for the deprecated Netapp parameters are converted to a valid Netapp driver configuration. As a result, Cinder's Netapp backend driver successfully initializes.

BZ#1590952

Previously, the Cinder Netapp backend ignored the CinderNetappNfsMountOptions TripleO Heat parameter that prevented configuration of the Netapp NFS mount options via the TripleO Heat parameter.
The code responsible for handling Cinder's Netapp configuration no longer ignores the CinderNetappNfsMountOptions parameter. The CinderNetappNfsMountOptions parameter correctly configures Cinder's Netapp NFS mount options.

BZ#1599409

During a version upgrade, Cinder's database synchronization is now executed only on the bootstrap node. This prevents database synchronization and upgrade failures that occurred when database synchronization was executed on all Controller nodes.