



Red Hat OpenStack Platform 13

Release Notes

Release details for Red Hat OpenStack Platform 13

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

Copyright © 2021 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

CHAPTER 1. INTRODUCTION	5
1.1. ABOUT THIS RELEASE	5
1.2. REQUIREMENTS	5
1.3. DEPLOYMENT LIMITS	5
1.4. DATABASE SIZE MANAGEMENT	6
1.5. CERTIFIED DRIVERS AND PLUG-INS	6
1.6. CERTIFIED GUEST OPERATING SYSTEMS	6
1.7. BARE METAL PROVISIONING SUPPORTED OPERATING SYSTEMS	6
1.8. HYPERVISOR SUPPORT	6
1.9. CONTENT DELIVERY NETWORK (CDN) REPOSITORIES	6
1.10. PRODUCT SUPPORT	8
1.11. UNSUPPORTED FEATURES	8
CHAPTER 2. TOP NEW FEATURES	9
2.1. RED HAT OPENSTACK PLATFORM DIRECTOR	9
2.2. CONTAINERS	9
2.3. BARE METAL SERVICE	9
2.4. CEPH STORAGE	9
2.5. COMPUTE	10
2.6. HIGH AVAILABILITY	10
2.7. METRICS AND MONITORING	11
2.8. NETWORK FUNCTIONS VIRTUALIZATION	12
2.9. OPENDAYLIGHT	12
2.10. OPENSTACK NETWORKING	12
2.11. SECURITY	12
2.12. STORAGE	13
2.13. TECHNOLOGY PREVIEWS	14
2.13.1. New Technology Previews	14
2.13.2. Previously Released Technology Previews	14
CHAPTER 3. RELEASE INFORMATION	17
3.1. RED HAT OPENSTACK PLATFORM 13 GA	17
3.1.1. Enhancements	17
3.1.2. Technology Preview	20
3.1.3. Release Notes	20
3.1.4. Known Issues	22
3.2. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - JULY 19, 2018	29
3.2.1. Enhancements	29
3.2.2. Release Notes	29
3.2.3. Known Issues	30
3.3. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - AUGUST 29, 2018	31
3.3.1. Enhancements	31
3.3.2. Release Notes	33
3.3.3. Known Issues	34
3.4. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - NOVEMBER 13, 2018	34
3.4.1. Enhancements	35
3.4.2. Release Notes	36
3.4.3. Known Issues	36
3.5. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - JANUARY 16, 2019	37
3.5.1. Enhancements	37
3.5.2. Known Issues	38

3.6. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - MARCH 13, 2019	38
3.6.1. Enhancements	38
3.6.2. Release Notes	40
3.6.3. Known Issues	40
3.6.4. Removed Functionality	40
3.7. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - APRIL 30, 2019	41
3.7.1. Enhancements	41
3.7.2. Known Issues	41
3.8. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - JULY 10, 2019	42
3.8.1. Enhancements	42
3.8.2. Technology Preview	42
3.8.3. Release Notes	43
3.8.4. Known Issues	43
3.9. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - SEPTEMBER 4, 2019	43
3.9.1. Enhancements	44
3.9.2. Technology Preview	45
3.9.3. Release Notes	45
3.9.4. Known Issues	46
3.9.5. Deprecated Functionality	46
3.10. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - NOVEMBER 6, 2019	46
3.10.1. Enhancements	46
3.11. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - DECEMBER 19, 2019	48
3.11.1. Enhancements	48
3.11.2. Deprecated Functionality	48
3.12. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - MARCH 10, 2020	48
3.12.1. Enhancements	48
3.13. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - JUNE 24, 2020	50
3.13.1. Bug Fix	50
3.13.2. Enhancements	51
3.13.3. Release Notes	51
3.14. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - OCTOBER 28, 2020	51
3.14.1. Bug Fix	52
3.14.2. Enhancements	53
3.14.3. Known Issues	54
3.15. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - DECEMBER 16, 2020	55
3.15.1. Bug Fix	55
3.16. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - MARCH 17, 2021	55
3.16.1. Bug Fix	55
3.16.2. Known Issues	55
3.17. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - JUNE 16, 2021	56
3.17.1. Bug Fix	56
3.17.2. Enhancements	58
3.17.3. Release Notes	58
CHAPTER 4. TECHNICAL NOTES	59
4.1. RHEA-2018:2086 – RED HAT OPENSTACK PLATFORM 13.0 ENHANCEMENT ADVISORY	59
4.2. RHSA-2018:2214 – IMPORTANT: OPENSTACK-TRIPLEO-HEAT-TEMPLATES SECURITY UPDATE	69
4.3. RHBA-2018:2215 – OPENSTACK-NEUTRON BUG FIX ADVISORY	71
4.4. RHBA-2018:2573 – OPENSTACK PLATFORM 13 BUG FIX AND ENHANCEMENT ADVISORY	72
4.5. RHBA-2018:2574 – OPENSTACK DIRECTOR BUG FIX ADVISORY	74
4.6. RHBA-2018:3587 – RED HAT OPENSTACK PLATFORM 13.0 DIRECTOR BUG FIX ADVISORY	80
4.7. RHBA-2019:0068 – RED HAT OPENSTACK PLATFORM 13 BUG FIX AND ENHANCEMENT ADVISORY	83
4.8. RHBA-2019:0448 – RED HAT OPENSTACK PLATFORM 13 BUG FIX AND ENHANCEMENT ADVISORY	85

4.9. RHBA-2021:2385 – RED HAT OPENSTACK PLATFORM 13 BUG FIX AND ENHANCEMENT ADVISORY 88

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 only the most recent release of Red Hat Enterprise Linux 7.

The Red Hat OpenStack Platform dashboard (horizon) 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

You can use Internet Explorer 11 to display the dashboard but expect a degradation of some functionalities because the browser is no longer maintained.

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"
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>.

1.11. UNSUPPORTED FEATURES

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

- Custom policies, which includes modification of **policy.json** files either manually or through any ***Policies** heat parameters. Do not modify the default policies unless the documentation contains explicit instructions to do so.

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

CHAPTER 2. TOP NEW FEATURES

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

2.1. 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. External clusters (those not deployed by director) running Red Hat

Ceph Storage 2.x (Jewel) will remain compatible with the newer Ceph client. 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
```

For the metrics to be created properly, ensure that 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, 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.

Removal of Telemetry API and ceilometer-collector

The Telemetry API service is replaced by the **OpenStack Telemetry Metrics** (gnocchi) service and the **OpenStack Telemetry Alarming** (aodh) service APIs. The **ceilometer-collector service** is replaced by the **ceilometer-notification-agent** daemon because the Telemetry polling agent sends the messages from the sample file to the **ceilometer-notification-agent** daemon.

**NOTE**

Ceilometer as a whole is not deprecated, just the Telemetry API service and the ceilometer-collector service.

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.

The following features are included:

- Date Plane Abstraction: A P4 plug-in for the platform.
- Containers: A plug-in for Kubernetes, as well as development of Neutron Northbound extensions for mixed VM-container environments.

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 the 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 known 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(629ff0024219488499b0fac0caciaa3a5). 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 workaroud 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 opendaylight_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 `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](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

HAProxy, used for network load balancing, must be version 1.6 or higher to correctly support the PING type health monitor.

The version of HAProxy included with Red Hat OpenStack Platform 13 is an older version than 1.6 that uses TCP connect instead when you configure the PING type health monitor.

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 nodes 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 pgscale (<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_pgs#set_the_number_of_

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
# iptables -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.

BZ#1790653

Because of the manner in which OpenStack Networking binds ports, the live migration of network instances in DVR environments might cause existing connections using a floating IP address to become disconnected. Currently, there is no workaround in RHOSP 13. However, this issue has been fixed in RHOSP 14 and later releases.

3.2. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - JULY 19, 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 OVSDDB 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 opendaylight_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 - AUGUST 29, 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(629ff0024219488499b0fac0caciaa3a5)`. 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`

3.4. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - NOVEMBER 13, 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.

3.4.1. Enhancements

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

BZ#1466117

To set MTU as a part of OSPD, this release adds **neutron::plugins::ml2::physical_network_mtus** as a NeutronML2PhysicalNetworkMtus in the heat template to enable MTU in the ml2 plugin. Neutron::plugins::ml2::physical_network_mtus is set based on values from the TripleO heat template.

BZ#1545151

Director uploads the latest amphora image to glance when OpenStack is updated and/or upgraded. The latest amphora image ensures amphora instances run with the latest general bug and security fixes, not only for Octavia agent fixes, but also for operating system fixes.

With this release, newly created and recreated amphora instances are made with the latest amphora image. Previous amphora images will remain stored in glance and be renamed to include the timestamp in the suffix.

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#1571286

You can use the weigher **CPUWeigher** to spread (default) or pack workloads on hosts based on their vCPU usage. You can set the nova.conf **[filter_scheduler] cpu_weight_multiplier** configuration option to -1.0 or 1.0. If you set this option to 1.0, instances are spread across hosts. If you set this option to -1.0, instances are packed on a host. If you set the value to 0, the weigher is disabled.

BZ#1619485

There is a case where **multipathd show status** doesn't return an error code as it should, so we are now checking stdout as a workaround for this issue to properly detect that multipathd is in an error state.

BZ#1628786

To prevent Ceph setup failures in a multi-architecture environment, this update ensures that Ceph setup of client users, keys, and pools is not run on a non-x86_64 system.

Ceph setup is supported only on x86_64 systems. Prior to this update, if the first inventory item in the clients group was not an x86_64 system, a Ceph setup attempt on that system would fail and cause the ceph-install to abort.

BZ#1629873

iSCSI device detection checked for the presence of devices based on the re-scan time. Devices becoming available between scans went undetected. With this release, searching and rescanning are independent operations working at different cadences with checks happening every second.

BZ#1631009

In prior versions, undercloud hieradata overrides could be used to tune some service configurations using the <service>::config options similar to the overcloud. However, this functionality was not available for all deployed OpenStack services. With this version, any configuration values not currently available can be updated via the <service>::config hieradata.

BZ#1640833

Support was added for volume retype and migration operations to the Block Storage service's HPE Nimble Storage driver.

3.4.2. 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#1496584

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

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

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 should be:

```
# 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#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 that causes a loss of the default routes.

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.

3.4.3. Known Issues

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

BZ#1621062

octavia-amphora file naming change can result in broken symlinks or symlinks with improper naming. As a workaround, run /usr/sbin/rhosp-director-image-update. This resolves the issue.

3.5. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - JANUARY 16, 2019

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

3.5.1. Enhancements

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

BZ#1476282

You can now deploy a minimal version of the overcloud with the **overcloud-minimal** qcow2 image. This minimal installation does not require OpenStack entitlements for updates.

BZ#1600115

Previously, the first packet of a new connection using an OVN logical router was used to discover the MAC address of the destination. This resulted in the loss of the first packet on the new connection.

This enhancement adds the capability to correctly queue the first packet of a new connection, which prevents the loss of that packet.

BZ#1607362

This feature adds support for deploying OpenDaylight (ODL) on IPv6 addresses.

BZ#1635892

Octavia previously assigned the Octavia project-id to the security group associated with the VIP and VRRP Amphora ports. This prevented the user from restricting access to the load-balancer. This fix adds the option to change SG ownership to belong to the user project (for certain whitelisted projects), which enables the user to refine access policies for the load-balancers.

BZ#1636395

This feature allows you to create instances with trusted SR-IOV virtual functions (VFs), such as changing the MAC address of the VF and enable promiscuous mode directly from the guest instance. These functions help you configure failover VFs for instances directly from the instance.

To configure trusted mode for VFs, you first set the **trusted** value of the **[pci] passthrough_whitelist** JSON configuration option in nova.conf. You then create the port with the **trusted=true** attribute in the binding profile. Make sure that the **vnic-type** attribute has the **hw_veb** or **direct** values.

BZ#1644020

This feature adds a new parameter **NovaLibvirtVolumeUseMultipath** (boolean), which sets the multipath configuration parameter **libvirt/volume_use_multipath** in the nova.conf file for Compute nodes. This parameter can be set for each Compute role. Default value is **False**.

BZ#1646907

This feature adds the capability to boot whole security-hardened images in UEFI mode.

BZ#1648261

This enhancement adds the parameter **NeutronOVSTunnelCsum**, which allows you to configure **neutron::agents::ml2::ovs::tunnel_csum** in the heat template. This parameter sets or removes the tunnel header checksum on the GRE/VXLAN tunnel that carries outgoing IP packets in the OVS agent.

BZ#1656495

This feature adds the parameter **NovaSchedulerWorkers**, which allows you to configure multiple **nova-schedule** workers for each scheduler node. Default value is **1**.

3.5.2. Known Issues

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

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.

1. Restart the faulty node.
2. 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](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#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
```

3.6. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - MARCH 13, 2019

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

3.6.1. Enhancements

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

BZ#1636496

With this update, you can use the following parameters to set the default Octavia timeouts for backend member and frontend client:

- **OctaviaTimeoutClientData:** Frontend client inactivity timeout
- **OctaviaTimeoutMemberConnect:** Backend member connection timeout
- **OctaviaTimeoutMemberData:** Backend member inactivity timeout
- **OctaviaTimeoutTcpInspect:** Time to wait for TCP packets for content inspection

The value for all of these parameters is in milliseconds.

BZ#1636895

With this update, you can now create tunnel endpoints on IPv6 addresses.

BZ#1650576

Previously, OpenDaylight packaging used the default OpenDaylight **log_pattern** values and included the PaxOsgi appender. These default values are not always appropriate for every deployment and it is appropriate to configure custom values.

With this update, **puppet-.opendaylight** has two additional configuration variables:

- **log_pattern:** Use this variable to configure which log pattern you want to use with the OpenDaylight logger log4j2.
- **enable_paxosgi_appender:** Use this boolean flag to enable or disable the PaxOsgi appender.

puppet-.opendaylight also modifies the OpenDaylight defaults. Deployments that use **puppet-.opendaylight** have new defaults:

- **log_pattern:** %d{ISO8601} | %-5p | %-16t | %-60c{6} | %m%n
- **enable_paxosgi_appender:** false

New variable configuration options

log_pattern

String that controls the log pattern used for logging.

Default: %d{ISO8601} | %-5p | %-16t | %-60c{6} | %m%n

Valid options: A string that is a valid log4j2 pattern.

enable_paxosgi_logger

Boolean that controls whether the PaxOsgi appender is enabled for logging.

If you enable the **enable_paxosgi_logger** variable, you must also modify the log pattern to utilize the additional capabilities. Modify the **log_pattern** variable and include a pattern that contains the PaxOsgi tokens. For example, set the **log_pattern** variable to a string that includes the following values:

```
'%X{bundle.id} - %X{bundle.name} -
+
%X{bundle.version}'
```

If you do not edit the **log_pattern** variable, the PaxOsgi appender is still enabled and continues to run but logging does not utilize the additional functionality.

For example, set the **enable_paxosgi_logger** variable to **true** and set the **log_pattern** variable to the following value:

```
'%d{ISO8601} | %-5p | %-16t | %-32c{1} | %X{bundle.id} - %X{bundle.name} - %X{bundle.version}
| %m%n'
```

Default: **false**

Valid options: The boolean values **true** and **false**.

3.6.2. 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#1597666

With this update, OpenDaylight minor update is now included in the Red Hat OpenStack Platform minor update workflow.

BZ#1611960

With this update, Compute nodes in a Red Hat OpenStack Platform environment that uses OpenDaylight as a back end can be scaled successfully.

3.6.3. Known Issues

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

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#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#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.

3.6.4. Removed Functionality

The following functionality has been removed from this release of Red Hat OpenStack Platform.

BZ#1431431

Block Storage - Highly Available Active-Active Volume Service is no longer available as a technology preview, and is not supported for this release.

BZ#1683464

Block Storage - RBD Cinder Volume Replication is no longer available for technology preview and is not supported.

3.7. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - APRIL 30, 2019

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

3.7.1. Enhancements

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

BZ#1654079

Previously, the undercloud upgrade failed if the overcloud was in a failed state. As a result, you could not deploy fixes by modifying the undercloud.

With this update, you can use the new undercloud upgrade option **--force**. This option causes the undercloud upgrade process to ignore the overcloud state.

BZ#1692872

With this update, you can use the following new tags to control the configuration management Ansible playbooks and tasks:

- container_config
- container_config_tasks
- container_config_scripts
- container_startup_configs
- host_config
- step1
- step2
- step3
- step4
- step5

3.7.2. Known Issues

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

BZ#1664701

A recent change made memory allocation for instances with NUMA topologies pagesize aware. With this change, memory for instances with NUMA topologies can no longer be oversubscribed.

As a result, memory oversubscription is currently disabled for all instances with a NUMA topology, whereas previously only instances with hugepages were not allowed to use oversubscription. This affects instances with an explicit NUMA topology and those with an implicit topology. An instance can have an implicit NUMA topology due to the use of hugepages or CPU pinning.

If possible, avoid the use of explicit NUMA topologies. If CPU pinning is required, resulting in an implicit NUMA topology, there is no workaround.

3.8. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - JULY 10, 2019

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

3.8.1. Enhancements

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

BZ#1589505

With this update you can now configure NUMA affinity to ensure instances are placed on the same NUMA node as the NIC providing external connectivity for the vSwitch. This feature is available for layer 2 networks of type 'flat' or 'vlan' and layer 3 networks of type 'vxlan', 'gre', or 'geneve'.

BZ#1688461

There are two new CLI arguments you can use with **config-download**:

- Monitor the deployment in a separate CLI session or with the API with **openstack overcloud status**.
- Log and save Ansible errors for future analysis with **openstack overcloud failures**.

BZ#1698467

This enhancement adds new features and usability enhancements to the Octavia Horizon dashboard.

BZ#1698683

This update adds a new director parameter, CinderNfsSnapshotSupport, which you can use to control Block Storage service (cinder) snapshots on NFS back ends.

BZ#1701427

Previously, if you enabled TLS throughout your environment, the communication between internal services, such as the haproxy and the manila API, was not secured.

With this update, the manila API supports TLS endpoints on the internal API network.

BZ#1714227

This update adds support for a second ceph Storage Tier deployment capability through director.

3.8.2. 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#1624971

With this Technology Preview update, you can attach a single volume to multiple hosts with the multi-attach feature. The volumes are connected without a storage protocol and data consistency must be provided by an appropriate software application such as a clustered filesystem.

Contact Red Hat to determine if your target storage driver supports the multi-attach feature in this release. Cinder volume encryption is not supported on multi-attached volumes and other feature restrictions may apply.

3.8.3. 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#1651191

This update adds support for the Ansible Networking ML2 plugin, which provides the following functionality:

- Isolated Tenant network capability for ironic baremetal guests.
- Automated switch configuration for baremetal nodes.
- Use of the same ML2 driver for multiple switch platforms.

BZ#1696332

This update provides OpenStack Messaging Service (zaqar) support for IPv6 endpoints, to facilitate an Undercloud deployed with IPv6.

BZ#1701425

Previously, the manila API was not deployed with Apache httpd server.

With this update, the Apache logs are located in `/var/log/containers/httpd/manila-api` on the nodes with manila API container.

3.8.4. Known Issues

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

BZ#1581337

HAProxy, used for network load balancing, must be version 1.6 or higher to correctly support the PING type health monitor.

The version of HAProxy included with Red Hat OpenStack Platform 13 is an older version than 1.6 that uses TCP connect instead when you configure the PING type health monitor.

3.9. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - SEPTEMBER 4, 2019

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

3.9.1. Enhancements

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

BZ#1614361

With this update, you can now use the Red Hat OpenStack Platform 13 host-config-and-reboot environment during fast-forward upgrade:

1. Remove the **NodeUserData** mapping from the **first-boot** script
2. Add the **host-config-and-reboot.yaml** environment file to the deploy command
3. Add KernelArgs and TunedProfile configured to the Ovs-DPDK role using role-specific parameters
4. Ensure that the KernelArgs and TunedProfile correspond to the OpenStack Platform 10 values. Any changes result in the node rebooting during fast-forward upgrade and the upgrade fails.

Ansible cannot handle reboots performed by the heat stack configuration. Any incorrect configuration that results in reboot causes the fast-forward upgrade process to fail.



NOTE

You can still perform the fast-forward upgrade with the existing first-boot scripts, even with the new patches present.

BZ#1631107

With this update, Red Hat OpenStack Platform contains a new parameter **DnsSearchDomains**. You can use this parameter for IDM and FreeIPA environments that have different DNS subdomains. Set this parameter in the **parameter_defaults** section of an environment file to add a list of DNS search domains to **resolv.conf**.

BZ#1679267

Previously, it was not possible to upgrade to TLS Everywhere in an existing deployment.

With this update, you can secure the in-flight connections between internal OpenStack services without reinstallation.

BZ#1688385

With this enhancement, you can pass arbitrary mysql config options to the mysql cluster on the overcloud with the **tripleo::profile::base::database::mysql::mysql_server_options** hiera hash key.

BZ#1706992

With this update, you can now configure automatic restart of instances on a Compute node if the node reboots without migrating the instances. You can configure Nova and the libvirt-guests agent to shut down the instances gracefully and restart them when the Compute node reboots. The following parameters are new in Red Hat OpenStack Platform 13: NovaResumeGuestsStateOnHostBoot (True/False) NovaResumeGuestsShutdownTimeout (default 300s)

BZ#1713761

With this update, you can set the domain for an ifcfg configuration with a key for interfaces called **domain**. Use this to improve DNS search, which is needed to support IDM/FreelIPA environments with different DNS subdomains.

BZ#1732220

With this update, rabbitmq-management interface is now enabled on localhost by default on the overcloud so that it is simpler to monitor and query the state of rabbitmq via its management API.

3.9.2. 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#1700882

This update expands the Block Storage service multi attach feature to the Ceph RBD driver.

3.9.3. 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#1592528

In rare circumstances, when you reboot a controller node several times, RabbitMQ can run in an inconsistent state that blocks 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(629ff0024219488499b0fac0caciaa3a5). Skipping it. - "openstack network agent list" returns that some agents are DOWN

To restore normal operation, run the following command on any one of the controller nodes: pcs resource restart rabbitmq-bundle

BZ#1656978

Previously, the Neutron SR-IOV agent set two possible states for virtual functions (VFs), **enable** or **disable**, which forced the VF link state regardless of the physical function (PF) link state.

With this update, the Neutron SR-IOV agent sets VFs to **auto** or **disable**. The **auto** state replicates the PF **up** or **down** automatically. As a result, if the PF is in the **down** state, the VF does not transmit or receive, even with other VFs in the same embedded switch (NIC).

**NOTE**

This behavior is not standard and depends on the NIC vendor implementation. Check the driver manual for the actual behavior of a VF in the **auto** state when the PF is **down**.

BZ#1708705

This update enables multi-attach capability for hpe3par driver.

3.9.4. Known Issues

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

BZ#1745130

There is currently a known issue for TLS Everywhere in-place upgrades, where overcloud nodes are consequently unable to enroll in IdM. As a workaround, remove `/etc/ipa/ca.crt/` from all overcloud nodes before running the overcloud deploy. For more information, see https://bugzilla.redhat.com/show_bug.cgi?id=1732564.

3.9.5. Deprecated Functionality

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

BZ#1541829

File injection from the Compute REST API. This will continue to be supported for now if using API microversion < 2.56. However, nova will eventually remove this functionality. The changes are as follows:

- Deprecate the **personality** parameter from the **POST /servers** create server API and the **POST /servers/{server_id}/action** rebuild server API. Specifying the **personality** parameter in the request body to either of these APIs will result in a **400 Bad Request** error response.
- Add support to pass **user_data** to the rebuild server API as a result of this change.
- Stop returning **maxPersonality** and **maxPersonalitySize** response values from the **GET /limits** API.
- Stop accepting and returning **injected_files**, **injected_file_path_bytes**, **injected_file_content_bytes** from the **os-quota-sets** and **os-quota-class-sets** APIs.

3.10. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - NOVEMBER 6, 2019

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

3.10.1. Enhancements

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

BZ#1561961

This release includes support for PCI device NUMA affinity policies. You can configure these policies as part of the **[pci]alias** configuration options. The following policies are supported:

- **required**
- **legacy**
- **preferred**

In all cases, strict NUMA affinity is provided if possible. The key difference between the policies is the amount of NUMA affinity you can sacrifice before failing to schedule.

You can use these policies to configure how strict your NUMA affinity is for each device, or more specifically, for each device alias. This is useful to ensure maximum resource utilization.

When you configure the 'preferred' policy for a PCI device, nova now utilizes CPUs on a different NUMA node from the NUMA node of the PCI device, if only a different node is available. This results in increased resource utilization with the downside of reduced performance for these instances.

BZ#1656292

With this update, you can now use NVIDIA vGPU GRID drivers with Red Hat OpenStack Platform 13. Red Hat fully supports installations that include these drivers, excluding support for third-party applications developed around the NVIDIA GPU and dedicated driver.

BZ#1684421

With this release, you can configure **NovaEnableRbdBackend** on a per-role basis so that you can deploy a subset of compute hosts with RBD ephemeral disks. The remaining hosts continue using the default local ephemeral disk.

NOTE For best performance, the images that you deploy to RBD ephemeral compute hosts must be in RAW format, and the images that you deploy to local ephemeral compute hosts must be in QCOW2 format.

BZ#1726733

Before this update, the default amphora timeouts were too high for production environments.

With this update, the default amphora timeouts are more suitable for production environments. You can also use new director parameters to override the defaults.

BZ#1731210

This enhancement upgrades facter to version 3, which improves performance when you deploy and run updates on systems with a large number of network interfaces. This version of facter supports fact caching and generates the fact list significantly faster.

NOTE You must run facter version 3 in the same containers that you deploy on the host system when you use the version of openstack-tripleo-heat-templates that implements facter version 3.

BZ#1732220

With this update, rabbitmq-management interface is now enabled on localhost by default on the overcloud so that it is simpler to monitor and query the state of rabbitmq with its management API.

BZ#1733260

With this update, **openstack-tripleo-common** is enhanced so that you can run the **openstack overcloud generate fencing** command to create proper fencing configuration for ironic nodes that use the **staging-ovirt** power driver, similar to virtual machines that you deploy on RHV.

puppet-tripleo has also been enhanced and now properly configures pacemaker **fence-rhevm** stonith agents for virtual machines on RHV.

BZ#1746112

With this update, the overcloud role name can now start with a number, for example, **10Controller** or **99Compute**.

BZ#1762167

This enhancement adds the following plugins to the collectd container: - connectivity - mysql - ping - procevent - snmp-agent - sysevent

3.11. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - DECEMBER 19, 2019

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

3.11.1. Enhancements

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

BZ#1766735

This enhancement allows OpenStack Image Service (glance) volumes to be mounted in a containerized control plane, as required by ScaleIO.

3.11.2. Deprecated Functionality

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

BZ#1719815

The OpenStack Telemetry Event Storage (panko) service is now deprecated. Support for panko is limited to usage from Red Hat Cloudforms only. Red Hat does not recommend using panko outside of the Red Hat Cloudforms use-case. You can use the following options instead of using panko:

- Poll the OpenStack Telemetry Metrics (gnocchi) service instead of polling panko. This gives you access to the resource history.
- Use the OpenStack Telemetry Alarming (aodh) service to trigger alarms when an event arises. You can use OpenStack Messaging Service (zaqar) to store alarms in a queue if an application cannot be reached directly by the OpenStack Telemetry Alarming (aodh) service.

3.12. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - MARCH 10, 2020

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

3.12.1. Enhancements

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

BZ#1726733

This enhancement provides default Octavia timeouts that are suitable for production environments, and the following new heat parameters to override the defaults:

- OctaviaConnectionMaxRetries
- OctaviaBuildActiveRetries
- OctaviaPortDetachTimeout

BZ#1757886

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

BZ#1760567

This update adds a feature that verifies that ceph-ansible is installed from the ceph-tools repository.

BZ#1760871

This update adds the following parameters to fine-tune Octavia keepalived VRRP settings:

octavia::controller::vrrp_advert_int

Amphora role and priority advertisement interval in seconds. Defaults to `::os_service_default`

octavia::controller::vrrp_check_interval

VRRP health check script run interval in seconds. Defaults to `::os_service_default`

octavia::controller::vrrp_fail_count

Number of successive failures before transition to a fail rate. Defaults to `::os_service_default`.

octavia::controller::vrrp_success_count

Number of consecutive successes before transition to a success rate. Defaults to `::os_service_default`.

octavia::controller::vrrp_garp_refresh_interval

Time in seconds between gratuitous ARP announcements from the MASTER. Defaults to `::os_service_default`.

octavia::controller::vrrp_garp_refresh_count

Number of gratuitous ARP announcements to make on each refresh interval. Defaults to `::os_service_default`.

To customize your environment, see https://access.redhat.com/documentation/en-us/red_hat_openshift_platform/13/html/advanced_overcloud_customization/.

BZ#1766735

This enhancement allows OpenStack Image Service (glance) volumes to be mounted in a containerized control plane, as required by ScaleIO.

BZ#1777993

With this update, support for the extension of attached volumes is now enabled for the Libvirt OpenStack Nova virt driver.

BZ#1782229

This enhancement adds two new parameters, **GlanceImageImportPlugins** and **GlanceImageConversionOutputFormat**, to enable Image service (glance) plugins during the image import process.

For example, if you enable the **image_conversion** plugin, the following command imports a qcow2 image, stores it in raw format, and converts it automatically after import:

```
glance image-create-via-import --disk-format qcow2 --container-format bare --name cirros --import-method web-download --uri http://download.cirros-cloud.net/0.4.0/cirros-0.4.0-x86\_64-disk.img
```

This means that you can store images always in raw format when you use RBD as an Image service driver.

3.13. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - JUNE 24, 2020

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

3.13.1. Bug Fix

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

BZ#1650046

Before this update, SELinux limitations prevented RabbitMQ logs from rotating correctly on the undercloud when logged in as root with **su**.

The **rabbitmq-server** package is updated and the **openstack-selinux** has a new policy rule to enable correct log rotation.

BZ#1783210

This update puts the correct `ipc:host` setting in the pacemaker template version of the cinder-backup container.

BZ#1805840

Before this update, deploying Ceph RadosGW did not create the **swiftoperator** role, and users with the **swiftoperator** role were not granted administration permissions for RadosGW Swift endpoints. With this update, deploying Swift or Ceph RadosGW creates the **swiftoperator** role automatically, and users with the **swiftoperator** role can now administer RadosGW objects as well as Swift objects.

BZ#1811122

Before this update, the xtremio driver reported incorrect available space capacity, which prevented virtual machine instances that rely on the storage backend from provisioning with insufficient space. After this update, the xtremio driver now reports the correct free space capacity, and virtual machine instances can provision correctly.

BZ#1813640

This update assigns higher priority to port create and update messages received by the neutron DHCP agent during processing to decrease 'Failed to allocate network' errors when booting instances.

BZ#1813642

This update fixes a bug that caused failure of upgrades from the OpenStack Platform maintenance release of 10 July 2019 (RHOSP 13.0.7). Specifically, it fixes a cell management error in OpenStack Compute (nova).

Now you can upgrade to newer OpenStack versions from the OpenStack Platform maintenance release of 10 July 2019 (RHOSP 13.0.7).

BZ#1829284

This update correctly associates a cinder volume with its corresponding VxFlex OS volume, so that when you delete a volume, the corresponding backend volume is also deleted. Previously, deletion of a cinder volume did not trigger deletion of the corresponding VxFlex OS volume.

BZ#1829765

Before this update, the cinder RBD driver did not perform trim or discard operations, which prevented users from trimming unused spaces from cinder RBD volumes.

With this update, the cinder RBD driver now supports trim and discard operations.

BZ#1835870

This update fixes a bug that caused Cinder offline volume migration for HPE 3par storage to fail.

3.13.2. Enhancements

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

BZ#1670592

This update adds support for hyper-covered infrastructure (HCI) deployments with OVS-DPDK. An HCI architecture features co-location of overcloud nodes with Compute and Ceph Storage services, configured for better resource usage.

BZ#1759254

This enhancement adds the **OctaviaConnectionLogging** parameter so that you can disable connection flow logging. The default setting is **true**, which means that connection flows are logged.

BZ#1823416

This enhancement updates the cinder Datera driver to version 2.2.

BZ#1841194

This enhancement corrects a condition that caused network traffic flooding with the openvswitch firewall driver. Previously, the lack of an forwarding database (FDB) entry for the destination MAC address of relevant packets caused flooding across all ports in a given VLAN on the integration bridge. Now the traffic is directed only to the correct port.

3.13.3. 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#1804412

This update lets you disable connection logging inside the amphora.

3.14. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - OCTOBER 28, 2020

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

3.14.1. Bug Fix

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

BZ#1723482

Before this update, the Compute (nova) service not releasing resources, such as network ports, until a Compute node is restored was causing the Load-balancing service (octavia) failover to fail when it was unable to detach a network port from an instance on a Compute node that is down.

With this update, the failover flow in the Load-balancing service has been updated to work around this Compute service issue. The Load-balancing service will now abandon ports that the Compute service will not release, leaving them in a "pending delete" state for the Compute service or Networking service to clean up once the Compute node is restored. This resolves the issue, allowing failover to succeed even if the Compute node is still failed.

BZ#1806975

Before this update, when several restores were being performed concurrently, the backup service was failing because the system was running out of memory.

With this update, we have increased the rate at which Python frees memory during backup restore operations by reducing the reference count to the data sooner, to allow Python to garbage collect the data as soon as it is decompressed rather than waiting until the restoration is complete. This resolves the issue, allowing the backup service to handle multiple restorations concurrently.

BZ#1841157

Before this update, FC live migration was failing. With this update, the correct device information is now sent to os-brick for FC for the corresponding host. Also, the device is now removed from the correct masking view when the live migration process has failed on the Compute node.

BZ#1841866

Before this update, the 3PAR driver did not examine the `_name_id` field for a possible volume ID, which caused volumes to be unusable after a live migration. With this update, the driver is now aware of the `_name_id` field as an alternative location for the volume ID, and live migrated volumes now work as expected.

BZ#1843196

Before this update, the internal temporary snapshot, created during async migration when creating a volume from a snapshot, was not being deleted from the VNX storage.

For example, if we create a new volume, V2, from snapshot S1, which we created from volume V1, an internal temporary snapshot, S2, is created from copying S1. V1 now has two snapshots, S1 and S2. Although we delete V1, V2 and S1 from OpenStack Block Storage (cinder), S2 is not deleted. This causes both V1 and S2 to remain on the VNX storage.

With this update, the temporary snapshot, S2, is deleted, and V1 can be successfully deleted.

BZ#1854950

Before this update, instances were unable to access their volumes after upgrading from RHOSP 10 to RHOSP 13, because the NFS share being used as a backend for OpenStack Block Storage (cinder) was not unmounted before migrating the OpenStack Block Storage services from the host to the containers. Therefore, when the containerized service started up and changed the ownership of all files in OpenStack Block Storage service directory, it also changed the ownership of files on the NFS share.

With this update, OpenStack Block Storage NFS shares are unmounted prior to upgrading the services to run in containers. This resolves the issue, and instances can now access their volumes after upgrading to RHOSP 13.

BZ#1861084

Before this update, if the OpenStack Shared File Systems (manila) service was configured with VServer-scoped ONTAP credentials it caused share provisioning to fail. This was due to a recent change to the NetApp ONTAP driver causing the share manager service to become stuck in a restart loop while trying to determine the storage system capabilities.

With this update, the NetApp ONTAP driver now checks for Vserver-scoped NetApp users and adds a fall-back path for determining storage system capabilities, which resolves the issue. The OpenStack Shared File Systems share manager service can now successfully determine the storage system capabilities, and share provisioning succeeds.

BZ#1862105

Before this update, initial connection errors with the agent were disrupting the retry logic, which sometimes resulted in the agent failing to communicate with the Ironic services, and logging a misleading `TypeError` to the agent console.

With this update, the exception handling has been fixed to explicitly handle known possible connection and lookup failure cases, and the logging has been updated to provide clarity on what is happening with the agent. Connections are now retried as designed by the agent and logging should no longer report just a `TypeError` in the event of an unexpected failure.

BZ#1867817

Before this update, using the **ceilometer-metrics-qdr.yaml** environment file resulted in a standalone Redis configuration rather than clustered redis instances as required by OpenStack Telemetry (ceilometer). This update now uses the correct services file in the resource registry to resolve the issue.

BZ#1847305

During startup, the **ironic-conductor** service could lose the reservation lock on a bare metal node for work accepted early in the **ironic-conductor** restart process. Losing the lock caused a race condition for work submitted to an OpenStack Bare Metal Provisioning (ironic) deployment during a restart of the **ironic-conductor** service, which resulted in requests failing with a "NodeNotLocked" error.

With this update, the database clean-up checks are performed before work can be accepted by an **ironic-conductor** process, which resolves the issue.

3.14.2. Enhancements

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

BZ#1802038

This enhancement adds support for an external Red Hat Ceph Storage 4.1 cluster.

BZ#1845802

This enhancement adds a new configuration option to the Networking (neutron) service, **http_retries**, which you can use to configure the number of times API calls to the Compute (nova) service or OpenStack Bare Metal Provisioning (ironic) service should be retried if the first attempt fails. By default, the API calls are retried 3 times in case of failure.

With this enhancement, the Networking service can retry API requests to prevent errors in booting instances, for example, to ensure that the Compute service receives notification that a port is ready for use.

BZ#1815202

When using the **config-download** Tech Preview functionality, the generated Ansible playbooks do not include a default **ansible.cfg** that is tailored to the **config-download** playbooks. The default Ansible settings are not ideal for large scale deployments.

This enhancement allows you to use the following command to generate an **ansible.cfg** that can be used with the **config-download** playbooks:

```
$ openstack tripleo config generate ansible
```

3.14.3. Known Issues

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

BZ#1891014

There is currently a known issue with TLS Everywhere environments when live migrating instances during a minor update.

With the introduction of support for full QEMU-native TLS encryption when live migrating (BZ1754791), instance live migration is failing when performing a minor update on a RHOSP deployment that has running instances. This is because the certificates for the TLS NBD block migration, that do not already exist in the libvirt container, are created during the update. The certificates are merged into the container directory tree during creation of the libvirt container, instead of being directly bind mounted from the host. Therefore, the QEMU processes of the instances that need migrated during the update do not get the new certificate automatically and the NBD setup process fails with the following error:

```
libvirtError: internal error: unable to execute QEMU command 'object-add': Unable to access credentials /etc/pki/qemu/ca-cert.pem: No such file or directory
```

Live migration works for instances created after the update.

Workaround:

You can use one of the following options to workaround this issue:

- Stop and start the instances that fail to live migrate after the update is complete, so that new QEMU processes get created by libvirt container that has the certificate details.
- Add the following configuration to the overcloud to disable TLS transport encryption for NBD, and deploy the overcloud:

```
parameter_defaults:
  UseTLSTransportForNbd: False
```

BZ#1726270

There is a known issue that causes the **glance db purge** command to fail with an **IntegrityError** and a message similar to "Cannot delete or update a parent row: a foreign key constraint fails" when you try to purge records from the images table when related records have not yet been purged from other tables.

Workaround:

Manually delete the records from other tables before you purge the records from the images table.

3.15. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - DECEMBER 16, 2020

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

3.15.1. Bug Fix

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

BZ#1879531

Before this update, when a Red Hat OpenStack Platform (RHOSP) user configured their containers to use the **latest** tag, RHOSP did not fetch nor rebuild these containers to use the updated container images.

With this update the issue is resolved. Now, anytime a user runs a deployment action (including an update), RHOSP always fetches the container images and checks the image ID for each running container to determine if it should be rebuilt to consume the latest image. RHOSP restarts any containers that it updates.



IMPORTANT

This update is a change from previous versions for how RHOSP manages container updates. In past versions, RHOSP would only check if the image existed. Now, RHOSP always refreshes containers during deployment actions, and restarts any containers that it updates. For this reason, you should not reuse tags, like **latest** and always use the **--tag-from-labels** option, unless you are controlling container tags with a Red Hat Satellite deployment.

3.16. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - MARCH 17, 2021

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

3.16.1. Bug Fix

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

BZ#1908366

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

Now the detachments do not fail.

3.16.2. Known Issues

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

BZ#1841371

When transferring a volume with snapshots to another user, the volume is transferred but the snapshots remain owned by the previous user.

As a workaround, manage snapshots manually in Red Hat Openstack Platform 13. See "Managing instance snapshots" [1] in the Red Hat OpenStack Instances and Images Guide.

[1] https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/13/html-single/instances_and_images_guide/index#section-instance-snapshots

3.17. RED HAT OPENSTACK PLATFORM 13 MAINTENANCE RELEASE - JUNE 16, 2021

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

3.17.1. Bug Fix

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

BZ#1888417

Before this update, API calls to the NetApp SolidFire back end for the Block Storage service (cinder) could fail with a **xNotPrimary** error. This type of error occurred when an operation was made to a volume at the same time that SolidFire automatically moved connections to rebalance the cluster workload.

With this update, a SolidFire driver patch adds the **xNotPrimary** exception to the list of exceptions that can be retried.

BZ#1888469

Before this update, users experienced timeouts in certain environments, mostly when volumes were too big. Often these multi-terabyte volumes experienced poor network performance or upgrade issues that involved the SolidFire cluster.

With this update, two timeout settings have been added to the SolidFire driver to allow users to set the appropriate timeouts for their environment.

BZ#1914590

Before this update, when a Block Storage service (cinder) API response was lost, the NetApp SolidFire back end created an unused duplicate volume.

With this update, a patch to the SolidFire driver first checks if the volume name already exists before trying to create it. The patch also checks for volume creation immediately after it detects a read timeout, and prevents invalid API calls.

BZ#1934440

Before this update, the Service Telemetry Framework (STF) client could not connect to the STF server, because the latest version of Red Hat AMQ Interconnect does not allow TLS connections without a CA certificate.

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

To obtain a Red Hat OpenShift TLS certificate, enter these commands:

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

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

```
MetricsQdrSSLProfiles:
- name: sslProfile
  caCertFileContent: |
    -----BEGIN CERTIFICATE-----
    ...
    TOpbgNIPcz0sloNK3Be0jUcYHVMPKGMR2kk=
    -----END CERTIFICATE-----
```

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

BZ#1940153

Before this update, when using the Block Storage service (cinder) to create a large number of instances (bootable volumes) from snapshots on HP3Par Storage back end servers, timeouts occurred. An HP variable (**convert_to_base**) was set to true which caused HP3Par to create a thick volume of the original volume. This was an unnecessary and unwanted action.

With this update, a newer HP driver (4.0.11) has been backported to RHOSP 13 that includes a new spec:

```
hpe3par:convert_to_base=True | False
```

- True (default) - The volume is created independently from the snapshot (HOS8 behavior).
- False - The volume is created as a child of snapshot (HOS5 behavior).

Usage

You can set this new spec for HPE3Par volumes by using the **cinder type-key** command:

```
cinder type-key <volume-type-name-or-ID> set hpe3par:convert_to_base=False | True
```

Example

```
$ cinder type-key myVolType set hpe3par:convert_to_base=False
$ cinder create --name v1 --volume-type myVolType 10
$ cinder snapshot-create --name s1 v1
$ cinder snapshot-list
$ cinder create --snapshot-id <snap_id> --volume-type myVolType --name v2 10
```

Notes

If the size of v2 is greater than size of v1, then the volume cannot be grown. In this case, to avoid any error, v2 is converted to a base volume (**convert_to_base=True**).

BZ#1943181

Before this update, when the Compute service (nova) made a **terminate-connection** call to the Block Storage service (cinder), single and multipath devices were not being flushed and there was a risk of data loss because these devices were in a **leftover** state.

The cause of this problem was that the os-brick **disconnect_volume** code assumed that the **use_multipath** parameter had the same value as the connector that was used in the original **connect_volume** call.

With this update, the Block Storage service changes how it performs disconnects. The os-brick code now properly flushes and detaches volumes when the multipath configuration in the Compute service changes for volumes attached to instances.

3.17.2. Enhancements

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

BZ#1875508

This enhancement enables you to override the Orchestration service (heat) parameter, **ServiceNetMap**, for a role when you deploy the overcloud.

On spine-and-leaf (edge) deployments that use TLS-everywhere, hiera interpolation has been problematic when used to map networks on roles. Overriding the ServiceNetMap per role fixes the issues seen in some TLS-everywhere deployments, provides an easier interface, and replaces the need for the more complex hiera interpolation.

3.17.3. 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#1924727

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

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 nodes 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 be 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

CollectdGnocchiKeystoneUserName

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 known 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 an 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 fixed 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. RHTSA-2018:2214 – IMPORTANT: OPENSTACK-TRIPLEO-HEAT-TEMPLATES SECURITY UPDATE

The bugs contained in this section are addressed by advisory RHTSA-2018:2214. Further information about this advisory is available at link: <https://access.redhat.com/errata/RHTSA-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 opendaylight_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 `UpgradelnitCommonCommand` 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 `UpgradelnitCommonCommand` 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-openshift

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 exceeded 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 nameservers 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 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 libvirt live-migration port range prevents live-migration failures. Previously, libvirt 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 libvirt 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.

4.6. RHBA-2018:3587 – RED HAT OPENSTACK PLATFORM 13.0 DIRECTOR BUG FIX ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2018:3587. Further information about this advisory is available at link: <https://access.redhat.com/errata/RHBA-2018:3587>

instack-undercloud

BZ#1627043

Some hardware changes boot device ordering in an unexpected way when receiving an IPMI bootdev command. This may prevent nodes from booting from the correct NIC or prevent PXE from booting at all. This release introduces a new "noop" management interface for the "ipmi" driver. When it is used, bootdev commands are not issued, and the current boot order is used. Nodes must be configured to try PXE booting from the correct NIC, and then fall back to the local hard drive. This change ensures a pre-configured boot order is kept with the new management interface.

BZ#1631009

In prior versions, undercloud hieradata overrides could be used to tune some service configurations using the <service>::config options similar to the overcloud. However, this functionality was not available for all deployed OpenStack services. With this version, any configuration values not currently available can be updated via the <service>::config hieradata.

openstack-tripleo-common

BZ#1631848

When upgrading from Red Hat OpenStack Platform 12 to 13 the ceph-osd package is removed. The

package removal stopped the running OSDs even though they were running in containers and shouldn't have required the package. This release removes the playbook that removes the package during the upgrade and Ceph OSDs are not unintentionally stopped during upgrade.

BZ#1545151

Director uploads the latest amphora image to glance when OpenStack is updated and/or upgraded. The latest amphora image ensures amphora instances run with the latest general bug and security fixes, not only for Octavia agent fixes, but also for operating system fixes.

With this release, newly created and recreated amphora instances are made with the latest amphora image. Previous amphora images will remain stored in glance and be renamed to include the timestamp in the suffix.

openstack-tripleo-heat-templates

BZ#1619092

One of the instance HA scripts connected to the publicURL keystone endpoint. This has now been moved to the internalURL endpoint by default. Additionally, an operator can override this via the '[placement]/valid_interfaces' configuration entry point in nova.conf.

BZ#1624899

In prior releases, triggers for online data migrations were missing. Online data migrations for nova, cinder, and ironic in the overcloud did not run automatically after upgrading to OSP 13, which forced a manual workaround. This release adds trigger logic for online data migrations. Online data migrations are triggered during the **openstack overcloud upgrade converge** command when upgrading to OSP 13.

BZ#1619104

In prior releases, you could set RX/TX queue size via `nova::compute::libvirt::rx_queue_size/nova::compute::libvirt::tx_queue_size`. However, there was no dedicated TripleO heat template parameter. With this release, the RX/TX queue size can be set on a role base like this:

```
parameter_defaults: ComputeParameters: NovaLibvirtRxQueueSize: 1024 NovaLibvirtTxQueueSize: 1024
```

The result is `rx_queue_size/tx_queue_size` is set using new parameters.

BZ#1466117

To set MTU as a part of OSPD, this release adds **neutron::plugins::ml2::physical_network_mtus** as a NeutronML2PhysicalNetworkMtus in the heat template to enable MTU in the ml2 plugin. `Neutron::plugins::ml2::physical_network_mtus` is set based on values from the TripleO heat template.

BZ#1594367

In prior versions, the conditions for checking whether a Docker daemon requires a restart were too strict. As a result, the Docker daemon and all containers were restarted whenever the Docker configuration changed or when the Docker RPM was updated. With this release, the conditions are relaxed to prevent

unnecessary container restarts. Use the "live restore" functionality for configuration changes to make sure the Docker daemon and all containers are restarted when Docker RPM is updated, but not when the Docker configuration is changed.

BZ#1612960

During a redeployment, a number of containers can be restarted needlessly, even in the absence of any configuration change. This was due to including too many unneeded files in the md5 calculation of the config files. With this release, no spurious container restarts are triggered by a redeploy.

BZ#1619663

The TripleO CinderNetappBackendName parameter did not correctly override the default value for cinder's Netapp back end. As a result, the name associated with cinder's Netapp back end could not be overridden. With this release, the CinderNetappBackendName parameter correctly overrides the default back end name.

puppet-cinder**BZ#1617199**

Several configuration settings were removed from cinder, but the corresponding parameters were not removed from the TripleO Puppet module responsible for setting cinder's configuration settings. As a result, invalid cinder configuration settings were added to cinder.conf. With this release, the Puppet module has been updated to prevent obsolete settings from being added to cinder.conf.

**NOTE**

The updated Puppet module will not remove any obsolete settings that were previously added to cinder.conf. Obsolete settings must be manually removed.

puppet-tripleo**BZ#1628705**

A faulty interaction between rhel-plugin-push.service and the Docker service occurred during system shutdown, which caused the controller reboot to take a long time. With this release, the correct shutdown ordering is enforced for these two services. Rebooting a controller takes less time now.

BZ#1602833

During deployment, an OVS switch may be configured with the incorrect OpenFlow controller port (6640, instead of 6653) for two out of the three controllers. This causes either a deployment failure, or a functional failure with the deployment later on, where the incorrect flows are programmed into the switch. This release correctly sets all of the OpenFlow controller ports to 6653 for each OVS switch. All of the OVS switches have the correct OpenFlow controller configuration, which consists of three URIs, one to each OpenDaylight using port 6653.

BZ#1488907

When a single OpenDaylight instance was removed from a cluster, this moved the instance into an isolated state, meaning it no longer acted on incoming requests. HA Proxy still load-balanced requests to the isolated OpenDaylight instance, which potentially resulted in OpenStack network commands

failing or not working correctly. HA Proxy now detects the isolated OpenDaylight instance as in an unhealthy state. HA Proxy does not forward requests to the isolated OpenDaylight.

python-os-brick

BZ#1631024

Under certain circumstances, the os-brick code responsible for scanning FibreChannel HBA hosts could return an invalid value. The invalid value would cause services such as cinder and nova to fail. With this release, the FibreChannel HBA scan code always returns a valid value. Cinder and nova no longer crash when scanning FibreChannel HBA hosts.

BZ#1607196

On multipath connections, devices are individually flushed for all paths upon disconnect. In certain cases, a failure on an individual device flush incorrectly prevents disconnection. With this release, individual paths are no longer flushed because flushing the multipath already ensures buffered data is written on the remote device. Now a disconnection only fails when it would actually lose data.

BZ#1619485

There is a case where **multipathd show status** doesn't return an error code as it should, so we are now checking stdout as a workaround for this issue to properly detect that multipathd is in an error state.

BZ#1628471

In prior releases, volume migration failed (with a VolumePathNotRemoved error) when a single iSCSI path failed in a narrow window a time during migration initiation. This release fixes the issue by extending the timeout period for verification of volume removal.

BZ#1629873

iSCSI device detection checked for the presence of devices based on the re-scan time. Devices becoming available between scans went undetected. With this release, searching and rescanning are independent operations working at different cadences with checks happening every second.

python-tripleoclient

BZ#1624462

In prior releases, if you used a custom plan—done via the '-p' option of the deploy command line—a number of passwords (such as mysql, horizon, pcsd, and so forth) were reset to new values during redeployment of an existing overcloud. This caused the redeployment to fail. With this release, a custom plan does not trigger setting new passwords.

4.7. RHBA-2019:0068 – RED HAT OPENSTACK PLATFORM 13 BUG FIX AND ENHANCEMENT ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2019:0068. Further information about this advisory is available at link: <https://access.redhat.com/errata/RHBA-2019:0068>

openstack-tripleo-common

BZ#1647931

Previously, when you updated a node from the undercloud, the **capabilities** field values were not always converted to a string value type. After this bug fix, the `capabilities` field now always converts to the string value type during node updates.

openstack-tripleo-heat-templates

BZ#1648261

This enhancement adds the parameter **NeutronOVSTunnelCsum**, which allows you to configure **neutron::agents::ml2::ovs::tunnel_csum** in the heat template. This parameter sets or removes the tunnel header checksum on the GRE/VXLAN tunnel that carries outgoing IP packets in the OVS agent.

BZ#1595114

OpenDaylight (ODL) configuration files were not recreated during Controller replacement, which caused updates to fail. This fix unmounts `/opt/.opendaylight/data` from the host, which causes the configuration files to be recreated during redeployment.

BZ#1641825

Previously, the OpenStack Platform Director did not configure authentication for Block Storage (Cinder) to access volumes that use the Nova privileged API. This caused operations on these volumes, such as migrating an in-use volume, to fail.

This bug fix adds the capability to configure Cinder with the Nova authentication data, which allows you to perform operations on volumes that use the privileged API with these credentials.

BZ#1652544

During an upgrade to a containerized deployment with Ironic, the TFTP server did not shut down correctly, which caused the upgrade to fail. This fix corrects the shutdown process of the TFTP server, so now the server can listen to the port and the upgrade completes successfully.

BZ#1638021

Previously, the inactivity probe timer for Open vSwitch from ODL was insufficient for larger-scale deployments, which caused the ODL L2 agents to appear as offline after the inactivity time lapsed.

This bug fix increases the default inactivity probe timer duration and adds the capability to configure the timer in the Director using the **OpenDaylightInactivityProbe** heat parameter. Default value is 180 seconds.

BZ#1639199

The location of Pacemaker log files for RabbitMQ containers was not set to the correct location, which caused unnecessary log files to be created in `/var/log/secure`. This fix adds mounting of the `/var/log/btmp` path during the start of the RabbitMQ container, which enables Pacemaker to create the logs in the correct location.

BZ#1644017

This feature adds the capability to configure the Cinder Dell EMC StorageCenter driver to use a multipath for volume-to-image and image-to-volume transfers. The feature includes a new parameter **CinderDellScMultipathXfer** with a default value of **True**. Enabling multipath transfers can reduce the total time of data transfers between volumes and images.

BZ#1644020

This feature adds a new parameter **NovaLibvirtVolumeUseMultipath** (boolean), which sets the multipath configuration parameter **libvirt/volume_use_multipath** in the nova.conf file for Compute nodes. This parameter can be set for each Compute role. Default value is **False**.

BZ#1656495

This feature adds the parameter **NovaSchedulerWorkers**, which allows you to configure multiple **nova-schedule** workers for each scheduler node. Default value is **1**.

BZ#1412661

Previously, the loopback device associated with an LVM volume group did not always restart after restarting the Controller node. This prevented the LVM volume groups used by the iSCSI Cinder backend from persisting after the restart, and prevented new volumes from being created.

After this bug fix, the loopback device is now restored after you reboot the Controller node, and the LVM volume group is accessible by the node.

BZ#1642446

This enhancement adds the **RabbitAdditionalErlArgs** parameter to the Erlang VM, which allows you to define custom arguments for the VM. The default argument is **+sbwt none**, which instructs the Erlang threads to go to sleep if no additional operations are required. For more information, see the Erlang documentation at: <http://erlang.org/doc/man/erl.html#+sbwt>

openstack-tripleo-heat-templates-compat**BZ#1593774**

Previously, OpenStack 13 Director did not set the correct Ceph version when deploying from OpenStack 12 templates. This caused the overcloud deployment to fail.

This bug fix sets the Ceph version to Jewel, and allows for correct deployment from OpenStack 12 templates.

puppet-opendaylight**BZ#1607362**

This feature adds support for deploying OpenDaylight (ODL) on IPv6 addresses.

4.8. RHBA-2019:0448 – RED HAT OPENSTACK PLATFORM 13 BUG FIX AND ENHANCEMENT ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2019:0448. Further information about this advisory is available at link: <https://access.redhat.com/errata/RHBA-2019:0448>

openstack-tripleo-common**BZ#1668774**

This bug was caused by updated versions of dmidecode 3.1 or later that returned system UUIDs in

lowercase. As a consequence, systems deployed with per node ceph-ansible customization prior to this version can break if UUID case mismatches and cause deployment failures. This fix updates the **openstack-tripleo-common** package to accept uppercase or lowercase UUIDs. Forced lowercase on dmidecode output make the code case insensitive.

openstack-tripleo-heat-templates

BZ#1659077

Previously, Octavia Health Manager did not receive heartbeat messages from amphorae due to a packet drop by the firewall. As a result, the **operating_status** of load balancers on Octavia composable role deployments never changed to **ONLINE**.

With this update, load balancers on Octavia composable role deployments change to **ONLINE** operating status successfully.

BZ#1636496

With this update, you can use the following parameters to set the default Octavia timeouts for backend member and frontend client:

- **OctaviaTimeoutClientData**: Frontend client inactivity timeout
- **OctaviaTimeoutMemberConnect**: Backend member connection timeout
- **OctaviaTimeoutMemberData**: Backend member inactivity timeout
- **OctaviaTimeoutTcpInspect**: Time to wait for TCP packets for content inspection

The value for all of these parameters is in milliseconds.

BZ#1655815

Previously, iSCSI connections that containerized OpenStack services created were not visible on the host. As a result, the host must close all iSCSI connections during shutdown. The shutdown sequence hung when the host failed to terminate these iSCSI connections and the host failed to terminate the hOpenStack connections because the connection information was not visible on the host.

With this update, connection information for containerized services that create iSCSI connections is now visible on the host, and the shutdown sequence no longer hangs.

BZ#1597666

With this update, OpenDaylight minor update is now included in the Red Hat OpenStack Platform minor update workflow.

BZ#1611960

With this update, Compute nodes in a Red Hat OpenStack Platform environment that uses OpenDaylight as a back end can be scaled successfully.

BZ#1623123

Previously, ODL configuration files were missing after redeployment.

With this update, **/opt/pendaylight/data** is no longer mounted on the host. As a result, the ODL configuration files are generated during redeployment.

BZ#1639203

Previously, the rabbitmq pacemaker bundle logged excessively during normal operation.

With this update, the rabbitmq bundle no longer logs excessively. In particular, the rabbitmq bundle does not log the harmless error **Failed to connect to system bus: No such file or directory**.

openstack-tripleo-image-elements

BZ#1646907

With this update, you can now boot whole security-hardened images in UEFI mode.

puppet-openshift

BZ#1650576

Previously, OpenDaylight packaging used the default OpenDaylight **log_pattern** values and included the PaxOsgi appender. These default values are not always appropriate for every deployment and it is appropriate to configure custom values.

With this update, **puppet-openshift** has two additional configuration variables:

- 1) **log_pattern**: Use this variable to configure which log pattern you want to use with the OpenDaylight logger log4j2.
- 2) **enable_paxosgi_appender**: Use this boolean flag to enable or disable the PaxOsgi appender.

puppet-openshift also modifies the OpenDaylight defaults. Deployments that use **puppet-openshift** have new defaults:

- **log_pattern**: %d{ISO8601} | %-5p | %-16t | %-60c{6} | %m%n
- **enable_paxosgi_appender**: false

New variable configuration options

log_pattern

String that controls the log pattern used for logging.

Default: %d{ISO8601} | %-5p | %-16t | %-60c{6} | %m%n

Valid options: A string that is a valid log4j2 pattern.

enable_paxosgi_logger

Boolean that controls whether the PaxOsgi appender is enabled for logging.

If you enable the **enable_paxosgi_logger** variable, you must also modify the log pattern to utilize the additional capabilities. Modify the **log_pattern** variable and include a pattern that contains the PaxOsgi tokens. For example, set the **log_pattern** variable to a string that includes the following values:

```
'%X{bundle.id} - %X{bundle.name} - %X{bundle.version}'
```

If you do not edit the **log_pattern** variable, the PaxOsgi appender is still enabled and continues to run but logging does not utilize the additional functionality.

For example, set the **enable_paxosgi_logger** variable to **true** and set the **log_pattern** variable to the following value:

```
'%d{ISO8601} | %-5p | %-16t | %-32c{1} | %X{bundle.id} - %X{bundle.name} - %X{bundle.version} | %m%n'
```

Default: **false**

Valid options: The boolean values **true** and **false**.

puppet-tripleo

BZ#1600449

Previously, deployments could fail when deploying the Overcloud with a BlockStorage role and setting a pacemaker property on nodes that belong to the BlockStorage role.

With this update, the pacemaker-managed cinder-volume resource starts only on nodes that pacemaker manages. As a result, Overcloud deployments with a BlockStorage role succeed.

4.9. RHBA-2021:2385 – RED HAT OPENSTACK PLATFORM 13 BUG FIX AND ENHANCEMENT ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2021:2385. Further information about this advisory is available at link: <https://access.redhat.com/errata/RHBA-2021:2385>

openstack-cinder component

BZ#1914590

Before this update, when a Block Storage service (cinder) API response was lost, the NetApp SolidFire back end created an unused duplicate volume.

With this update, a patch to the SolidFire driver first checks if the volume name already exists before trying to create it. The patch also checks for volume creation immediately after it detects a read timeout, and prevents invalid API calls. (BZ#1914590)

BZ#1940153

Before this update, when using the Block Storage service (cinder) to create a large number of instances (bootable volumes) from snapshots on HP3Par Storage back end servers, timeouts occurred. An HP variable (**convert_to_base**) was set to true which caused HP3Par to create a thick volume of the original volume. This was an unnecessary and unwanted action.

With this update, a newer HP driver (4.0.11) has been backported to RHOSP 13 that includes a new spec:

```
hpe3par:convert_to_base=True | False
```

- True (default) - The volume is created independently from the snapshot (HOS8 behavior).
- False - The volume is created as a child of snapshot (HOS5 behavior).

Usage

You can set this new spec for HPE3Par volumes by using the **cinder type-key** command:

```
cinder type-key <volume-type-name-or-ID> set hpe3par:convert_to_base=False | True
```

Example

```
$ cinder type-key myVolType set hpe3par:convert_to_base=False
$ cinder create --name v1 --volume-type myVolType 10
$ cinder snapshot-create --name s1 v1
$ cinder snapshot-list
$ cinder create --snapshot-id <snap_id> --volume-type myVolType --name v2 10
```

Notes

If the size of v2 is greater than size of v1, then the volume cannot be grown. In this case, to avoid any error, v2 is converted to a base volume (**convert_to_base=True**). (BZ#1940153)

BZ#1888417

Before this update, API calls to the NetApp SolidFire back end for the Block Storage service (cinder) could fail with a **xNotPrimary** error. This type of error occurred when an operation was made to a volume at the same time that SolidFire automatically moved connections to rebalance the cluster workload.

With this update, a SolidFire driver patch adds the **xNotPrimary** exception to the list of exceptions that can be retried. (BZ#1888417)

BZ#1888469

Before this update, users experienced timeouts in certain environments, mostly when volumes were too big. Often these multi-terabyte volumes experienced poor network performance or upgrade issues that involved the SolidFire cluster.

With this update, two timeout settings have been added to the SolidFire driver to allow users to set the appropriate timeouts for their environment. (BZ#1888469)

openstack-tripleo-heat-templates

BZ#1875508

This enhancement enables you to override the Orchestration service (heat) parameter, **ServiceNetMap**, for a role when you deploy the overcloud.

On spine-and-leaf (edge) deployments that use TLS-everywhere, hiera interpolation has been problematic when used to map networks on roles. Overriding the ServiceNetMap per role fixes the issues seen in some TLS-everywhere deployments, provides an easier interface, and replaces the need for the more complex hiera interpolation. (BZ#1875508)

BZ#1924727

The Block Storage backup service sometimes needs access to files on the host that would otherwise not be available in the container that runs the service. This enhancement adds the **CinderBackupOptVolumes** parameter, which you can use to specify additional container volume mounts for the Block Storage backup service. (BZ#1924727)

puppet-tripleo

BZ#1934440

Before this update, the Service Telemetry Framework (STF) client could not connect to the STF server, because the latest version of Red Hat AMQ Interconnect does not allow TLS connections without a CA certificate.

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

To obtain a Red Hat OpenShift TLS certificate, enter these commands:

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

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

```
MetricsQdrSSLProfiles:
- name: sslProfile
  caCertFileContent: |
    -----BEGIN CERTIFICATE-----
    ...
    TOpbgNIPcz0sloNK3Be0jUcYHVMPKGMR2kk=
    -----END CERTIFICATE-----
```

Then, redeploy your overcloud with the **openstack overcloud deploy** command. (BZ#1934440)

python-os-brick

BZ#1943181

Before this update, when the Compute service (nova) made a **terminate-connection** call to the Block Storage service (cinder), single and multipath devices were not being flushed and there was a risk of data loss because these devices were in a **leftover** state.

The cause of this problem was that the os-brick **disconnect_volume** code assumed that the **use_multipath** parameter had the same value as the connector that was used in the original **connect_volume** call.

With this update, the Block Storage service changes how it performs disconnects. The os-brick code now properly flushes and detaches volumes when the multipath configuration in the Compute service changes for volumes attached to instances. (BZ#1943181)