



Red Hat OpenStack Platform 10

Release Notes

Release details for Red Hat OpenStack Platform 10

Red Hat OpenStack Platform 10 Release Notes

Release details for Red Hat OpenStack Platform 10

OpenStack Documentation Team
Red Hat Customer Content Services
rhos-docs@redhat.com

Legal Notice

Copyright © 2016 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](#). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

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, 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 Software Collections 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	3
1.1. ABOUT THIS RELEASE	3
1.2. REQUIREMENTS	3
1.3. DEPLOYMENT LIMITS	4
1.4. DATABASE SIZE MANAGEMENT	4
1.5. CERTIFIED DRIVERS AND PLUG-INS	4
1.6. CERTIFIED GUEST OPERATING SYSTEMS	4
1.7. BARE METAL PROVISIONING SUPPORTED OPERATING SYSTEMS	4
1.8. HYPERVISOR SUPPORT	4
1.9. CONTENT DELIVERY NETWORK (CDN) CHANNELS	5
1.10. PRODUCT SUPPORT	6
CHAPTER 2. TOP NEW FEATURES	8
2.1. RED HAT OPENSTACK PLATFORM DIRECTOR	8
2.2. COMPUTE	8
2.3. DASHBOARD	9
2.4. IDENTITY	9
2.5. OBJECT STORAGE	10
2.6. OPENSTACK NETWORKING	10
2.7. SHARED FILE SYSTEM	11
2.8. TELEMETRY	11
2.9. HIGH AVAILABILITY	11
2.10. BARE METAL PROVISIONING SERVICE	12
2.11. OPENSTACK INTEGRATION TEST SUITE SERVICE	12
2.12. OPENSTACK DATA PROCESSING SERVICE	13
2.13. TECHNOLOGY PREVIEWS	13
CHAPTER 3. RELEASE INFORMATION	16
3.1. RED HAT OPENSTACK PLATFORM 10 GA	16
3.2. RED HAT OPENSTACK PLATFORM 10 MAINTENANCE RELEASES - MAY 2018 UPDATE	30
3.3. RED HAT OPENSTACK PLATFORM 10 MAINTENANCE RELEASE 26 JUNE 2018	39
3.4. RED HAT OPENSTACK PLATFORM 10 MAINTENANCE RELEASE 17 SEPTEMBER 2018	42
3.5. RED HAT OPENSTACK PLATFORM 10 MAINTENANCE RELEASE 16 JANUARY 2019	45
CHAPTER 4. TECHNICAL NOTES	47
4.1. RHEA-2016:2948 — RED HAT OPENSTACK PLATFORM 10 ENHANCEMENT UPDATE	47
4.2. RHEA-2018:2670 — RED HAT OPENSTACK PLATFORM 10 ENHANCEMENT UPDATE	69
4.3. RHEA-2018:2671 — RED HAT OPENSTACK PLATFORM 10 ENHANCEMENT UPDATE	74
4.4. RHBA-2019:0055 — RED HAT OPENSTACK PLATFORM 10 BUG FIX UPDATE	77

CHAPTER 1. INTRODUCTION

Red Hat OpenStack Platform provides the foundation to build a private or public Infrastructure-as-a-Service (IaaS) cloud on Red Hat Enterprise Linux. It offers a massively scalable, fault-tolerant platform for the development of cloud-enabled workloads.

The current Red Hat system is based on OpenStack Newton, and packaged so that available physical hardware can be turned into a private, public, or hybrid cloud platform including:

- Fully distributed object storage
- Persistent block-level storage
- Virtual-machine provisioning engine and image storage
- Authentication and authorization mechanism
- Integrated networking
- Web browser-based GUI for both users and administration.

The Red Hat OpenStack Platform IaaS cloud is implemented by a collection of interacting services that control its computing, storage, and networking resources. The cloud is managed using a web-based interface which allows administrators to control, provision, and automate OpenStack resources. Additionally, the OpenStack infrastructure is facilitated through an extensive API, which is also available to end users of the cloud.

1.1. ABOUT THIS RELEASE

This release of Red Hat OpenStack Platform is based on the OpenStack "Newton" 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 "Newton" release itself are available at the following location:

<https://releases.openstack.org/newton/index.html>

Red Hat OpenStack Platform uses components from other Red Hat products. See the following links for specific information pertaining to the support of these components:

<https://access.redhat.com/site/support/policy/updates/openstack/platform/>

To evaluate Red Hat OpenStack Platform, sign up at:

<http://www.redhat.com/openstack/>.



NOTE

The Red Hat Enterprise Linux High Availability Add-On is available for Red Hat OpenStack Platform use cases. See the following URL for more details on the add-on: <http://www.redhat.com/products/enterprise-linux-add-ons/high-availability/>. See the following URL for details on the package versions to use in combination with Red Hat OpenStack Platform: <https://access.redhat.com/site/solutions/509783>

1.2. REQUIREMENTS

Red Hat OpenStack Platform supports the most recent release of Red Hat Enterprise Linux. This version of Red Hat OpenStack Platform is supported on Red Hat Enterprise Linux 7.3.

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

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



NOTE

Prior to deploying Red Hat OpenStack Platform, it is important to consider the characteristics of the available deployment methods. For more information, refer to the [Installing and Managing Red Hat OpenStack Platform](#).

1.3. DEPLOYMENT LIMITS

For a list of deployment limits for Red Hat OpenStack Platform, see [Deployment Limits for Red Hat OpenStack Platform](#).

1.4. DATABASE SIZE MANAGEMENT

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

1.5. CERTIFIED DRIVERS AND PLUG-INS

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

1.6. CERTIFIED GUEST OPERATING SYSTEMS

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

1.7. BARE METAL PROVISIONING SUPPORTED OPERATING SYSTEMS

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

1.8. HYPERVISOR SUPPORT

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

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

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

1.9. CONTENT DELIVERY NETWORK (CDN) CHANNELS

This section describes the channel and repository settings required to deploy Red Hat OpenStack Platform 10.

You can install Red Hat OpenStack Platform 10 through the Content Delivery Network (CDN). To do so, configure **subscription-manager** to use the correct channels.



WARNING

Do not upgrade to the Red Hat Enterprise Linux 7.3 kernel without also upgrading from Open vSwitch (OVS) 2.4.0 to OVS 2.5.0. If only the kernel is upgraded, then OVS will stop functioning.

Run the following command to enable a CDN channel:

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

Run the following command to disable a CDN channel:

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

Table 1.1. Required Channels

Channel	Repository Name
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 10 for RHEL 7 (RPMs)	rhel-7-server-openstack-10-rpms
Red Hat Enterprise Linux 7 Server - Extras (RPMs)	rhel-7-server-extras-rpms

Table 1.2. Optional Channels

Channel	Repository Name
Red Hat Enterprise Linux 7 Server - Optional	<code>rhel-7-server-optional-rpms</code>
Red Hat OpenStack Platform 10 Operational Tools for RHEL 7 (RPMs)	<code>rhel-7-server-openstack-10-optools-rpms</code>

Channels to Disable

The following table outlines the channels you must disable to ensure Red Hat OpenStack Platform 10 functions correctly.

Table 1.3. Channels to Disable

Channel	Repository Name
Red Hat CloudForms Management Engine	<code>"cf-me-"</code>
Red Hat Enterprise Virtualization	<code>"rhel-7-server-rhev"</code>
Red Hat Enterprise Linux 7 Server - Extended Update Support	<code>"*-eus-rpms"</code>

**WARNING**

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

1.10. PRODUCT SUPPORT

Available resources include:

Customer Portal

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

- Knowledge base articles and solutions.
- Technical briefs.
- Product documentation.

- Support case management.

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

Mailing Lists

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

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

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

CHAPTER 2. TOP NEW FEATURES

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

2.1. RED HAT OPENSTACK PLATFORM DIRECTOR

This section outlines the top new features for the director.

Custom Roles and Composable Services

Monolithic templates have been decomposed into a set of multiple smaller discrete templates, each representing a composable service. These can be deployed on a standalone node, or combined with other services in the form of Custom Roles.

Note the following guidelines and limitations for the composable node architecture:

- You can assign any **systemd** managed service to a supported standalone custom role.
- You cannot split Pacemaker-managed services. This is because the Pacemaker manages the same set of services on each node within the overcloud cluster. Splitting Pacemaker-managed services can cause cluster deployment errors. These services should remain on the Controller role.
- You cannot change to custom roles and composable services during the upgrade process from Red Hat OpenStack Platform 9 to 10. The upgrade scripts can only accommodate the default overcloud roles.
- You can create additional custom roles after the initial deployment and deploy them to scale existing services.
- You cannot modify the list of services for any role after deploying an overcloud. Modifying the service lists after Overcloud deployment can cause deployment errors and leave orphaned services on nodes.

For more information on supported architecture for custom roles and composable services, see [Composable Services and Custom Roles](#) in the *Advanced Overcloud Customization* guide.

Graphical User Interface

Director can now be managed using a Graphical User Interface, which includes integrated templates, a built-in workflow, and pre- and post-flight validation checking. You use the GUI to create Role Assignments and perform node registration and introspection.

Separation of the Hardware Deployment Phase and Generic Node Deployment

The director workflow now includes a clear separation of the hardware deployment phase. This delineates where a user registers hardware to the inventory, uploads images, and defines hardware profiles. This phase is completed by deploying a given image to a specific hardware node. This separation allows you to deploy Red Hat Enterprise Linux onto a hardware node and hand it over to a user.

2.2. COMPUTE

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

Guest Device Role Tagging and Metadata Injection

With this update, OpenStack Compute creates and injects an additional metadata file which allows the guest to identify the instance based on the tags - the type of device, the bus it is attached to, the device address, the MAC address or drive serial string, the network or disk device name.

The guest is allowed to interpret the data. When the device role tags are used, the data is available through the metadata server and the configuration drive. For example, an example metadata file is as follows:

```
{ "devices": [ {
    "type": "nic",
    "bus": "pci",
    "address": "0000:00:02.0",
    "mac": "01:22:22:42:22:21",
    "tags": ["nfvfunc1"]
  }, {
    "type": "disk",
    "bus": "scsi",
    "address": "1:0:2:0",
    "serial": "disk-vol-2352423",
    "tags": ["dbvolume"]
  }
]
```

Newly defined API policy defaults

The API policy defaults are now defined in code like configuration options. Because of this, the sample `policy.json` file that is shipped with the Compute service (nova) is empty and should only be necessary if you want to override the API policy from the defaults in the code.

To generate the policy file you can run:

```
# oslopolicy-sample-generator --config-file=/etc/nova/nova-policy-generator.conf
```

2.3. DASHBOARD

This section outlines the top new features for the Dashboard.

Improved User Experience

The Swift panel is now rendered in AngularJS. This provides a hierarchy view of stored objects, client-side pagination, search, sorting of objects stored in Swift.

In addition, this release adds support for multiple, dynamically-set themes.

Improved Parity with Core OpenStack Services

This release now supports domain-scoped tokens (required for identity management in Keystone V3). Also, this release adds support for launching Nova instances attached to an SR-IOV port.

2.4. IDENTITY

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

Fernet Token Support

Red Hat OpenStack Platform 10 adds Fernet token support. The lightweight Fernet tokens mean that only minimal identity information is required. The non-persistent state means that no database backend is needed. Symmetric encryption has been implemented using AES-CBC signed with SHA256HMAC. As a result, you can expect significant performance improvement over UUID tokens.

Multi-domain LDAP Support

This release adds director support for multi-domain LDAP integration, allowing you to use multiple back ends for user authentication.

Expanded Role Capabilities

Red Hat OpenStack Platform 10 has expanded the role capabilities with **Domain-specific roles** and **Implied Roles**. Domain-specific roles - Allow role definition to be limited to a specific domain. These roles can be then assigned to a domain or project within the domain. Implied Roles - Inference rules can state that assignment of one role implies the assignment of another. These changes are expected to make role management much easier for administrators.

2.5. OBJECT STORAGE

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

Update Container on Fast-POST

This feature allows fast, efficient updates of metadata without the need to fully re-copy the contents of an object.

2.6. OPENSTACK NETWORKING

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

Full support for Distributed Virtual Routing

DVR is now fully supported in Red Hat OpenStack Platform 10. Users are able to choose between centralized routing (the default), and DVR. With DVR, each Compute node manages routing functionality.

Users are advised to refer to the documentation, and carefully plan whether centralized routing or DVR better suits their needs and overall network architecture.

DSCP Markings

Open vSwitch can now add DSCP marks to outbound network traffic, as defined in RFC 2474.

Enhanced NFV Datapath with Director Integration

Red Hat OpenStack Platform 10 adds support for SR-IOV PF passthrough (using **vnic_type=direct-physical**), in addition to VF passthrough. SR-IOV deployment can now be automated using the director. In addition, OVS-DPDK 2.5 is now fully supported and integrated with director.



WARNING

Do not upgrade to the Red Hat Enterprise Linux 7.3 kernel without also upgrading from Open vSwitch (OVS) 2.4.0 to OVS 2.5.0. If only the kernel is upgraded, then OVS will stop functioning.

Networking VLAN Aware Virtual Machines

Certain types of virtual machines require the ability to pass VLAN-tagged traffic over one interface, which is now represented as a `trunk` networking port. To create a trunk for use by a virtual machine, users must create a single parent port and one or more sub-ports. All of the ports and respective networks will be available to the interface, which should tag traffic on its interface using 802.1q.

2.7. SHARED FILE SYSTEM

This section outlines the top new features for the Shared File System service.

Director Integration

The Shared File System service (manila) is now a composable controller service deployable through director, and is now fully supported. With this release, the NetApp driver is also fully integrated into director, thereby enabling NetApp back end configuration for the Shared File System service out-of-the-box. The CephFS native driver (Technology Preview) is also fully integrated into director.

2.8. TELEMETRY

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

New Telemetry Meter Dispatcher Backend: Gnocchi

The Telemetry service (ceilometer) now uses **gnocchi** as its default meter dispatcher backend. Gnocchi is more scalable and is more aligned to the future direction of the Telemetry service. In addition, the Gnocchi backend also disables the legacy Ceilometer API in favor of the newer Gnocchi API.

2.9. HIGH AVAILABILITY

This section outlines the top new features for high availability.

Updated Service Management

The majority of core OpenStack services, including memcached and others, are now managed by systemd. A minimal number of critical services remain under Pacemaker, with fencing where required: HAProxy/virtual IPs, RabbitMQ, Galera (MariaDB), Manila-share, Cinder Volume, Cinder Backup, Redis.

Operational and Monitoring Tools

Red Hat OpenStack Platform 10 includes support for exposing information to operational and monitoring tools for High Availability.

2.10. BARE METAL PROVISIONING SERVICE

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

Standard Bare Metal to Tenant Support

With this release, the Bare Metal Provisioning service adds tenant support for the overcloud. This feature allows for a pool of shared hardware resources to be provisioned on demand by the cloud tenants.

Bare Metal Provisioning Certification Program

This release introduces the Bare Metal Provisioning driver certification program. This program provides assurance of hardware lifecycle management for both the infrastructure and the bare metal to tenant use cases.

2.11. OPENSTACK INTEGRATION TEST SUITE SERVICE

This section outlines the top new features for the OpenStack Integration Test Suite (tempest) service.

Overall Tempest Cleanup

This update includes an overall **tempest** cleanup including remote client debuggability, documentation review, client and manager aliases, and refactored test base class setup and teardown steps.

Refactored Tempest CLI

This update adds a domain-specific **tempest run** command that can be used as the primary entry point for running **tempest** tests.

Updates to Negative Test Guidelines

While the existing negative tests remain, this update adds support to the negative tests at the component level.

Migrated Python Repository

With this update, the **tempest-lib** Python repository is now migrated to the **tempest/lib** directory in the **tempest** repository.

Client Manager Refactor

Previously, the client managers instantiated all available clients at `__init__` time regardless of the `tempest` configuration about the available services, extensions and API versions and exposed the clients using the class attributes. With this release, clients are only instantiated on demand, and the manager internally caches instances of clients and serves them from the cache where applicable.`

Test Resource Management

With this release, all test resources are managed in a dedicated YAML file, which allows for the **tempest** configuration to happen with the same amount of configuration a deployer system uses to configure the OpenStack services. This also ensures that the test can select resources to be used by

the logical name or properties (for example, use whatever image that would fit in the 'smallest' flavor), or run against all combinations of certain resources.

Microversion Tests

This release adds some Compute Microversion tests to the Microversion testing framework.

2.12. OPENSTACK DATA PROCESSING SERVICE

This section outlines the top new features for the OpenStack Data Processing (sahara) service.

Support for the Latest Versions of the Most Popular Big Data Platforms and Components

This release adds support for Hortonworks Data Platform 2.3, 2.4 Stack and the new MapR 5.1 plugin (add-mapr-510).

Improved User Experience and Ease of Use

This release adds CLI for plugin-declared image creation, by enabling plugins to specify **yaml**-based recipes for image packing. It also include new CLI tools that enable users to easily generate image based on specification.

This release also adds integration with the Dashboard using the **openstack-sahara-ui** package.

2.13. TECHNOLOGY PREVIEWS

This section outlines features that are in technology preview in Red Hat OpenStack Platform 10.



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:

At-Rest Encryption

Objects can now be stored in encrypted form (using AES in CTR mode with 256-bit keys). This provides options for protecting objects and maintaining security compliance in Object Storage clusters.

Erasure Coding (EC)

The Object Storage service includes an EC storage policy type for devices with massive amounts of data that are infrequently accessed. The EC storage policy uses its own ring and configurable set of parameters designed to maintain data availability while reducing cost and storage requirements (by requiring about half of the capacity of triple-replication). Because EC requires more CPU and network resources, implementing EC as a policy allows you to isolate all the storage devices associated with your cluster's EC capability.

Open vSwitch Firewall Driver

The OVS firewall driver is now available as a Technology Preview. The conntrack-based firewall driver can be used to implement Security Groups. With conntrack, Compute instances are connected directly to the integration bridge for a more simplified architecture and improved performance.

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:

1. 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
2. Deploy Engines - deploy an OpenStack distribution before any benchmarking procedures take place, using servers retrieved from Server Providers
3. Verification - runs specific set of tests against the deployed cloud to check that it works correctly, collects results & presents them in human readable form
4. Benchmark Engine - allows to write parameterized benchmark scenarios & run them against the cloud.

Cells

OpenStack Compute includes the concept of Cells, provided by the nova-cells package, for dividing computing resources. For more information about Cells, see [Schedule Hosts and Cells](#).

Alternatively, Red Hat OpenStack Platform also provides fully supported methods for dividing compute resources in Red Hat OpenStack Platform; namely, Regions, Availability Zones, and Host Aggregates. For more information, see [Manage Host Aggregates](#).

CephFS Native Driver for Manila

The CephFS native driver allows the Shared File System service to export shared CephFS file systems to guests through the Ceph network protocol. Instances must have a Ceph client installed to mount the file system. The CephFS file system is included in Red Hat Ceph Storage 2.0 as a technology preview as well.

Containerized Compute Nodes

The Red Hat OpenStack Platform director has the ability to integrate services from OpenStack's containerization project (kolla) into the Overcloud's Compute nodes. This includes creating Compute nodes that use Red Hat Enterprise Linux Atomic Host as a base operating system and individual containers to run different OpenStack services.

DNS-as-a-Service (DNSaaS)

Red Hat OpenStack Platform 8 includes a Technology Preview of DNS-as-a-Service (DNSaaS), also known as Designate. DNSaaS 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. In addition, DNSaaS includes integration support for PowerDNS and Bind9.

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 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.

OpenDaylight Integration

Red Hat OpenStack Platform 10 includes a technology preview of integration with the OpenDaylight SDN controller. OpenDaylight is a flexible, modular, and open SDN platform that supports many different applications. The OpenDaylight distribution included with Red Hat OpenStack Platform 10 is limited to the modules required to support OpenStack deployments using NetVirt, and is based on the upstream Boron version.

The following packages provide the Technology Preview: `opendaylight`, `networking-odl`.

For more information, see the [Red Hat OpenDaylight Product Guide](#) and the [OpenDaylight and Red Hat OpenStack Installation and Configuration Guide](#).

Real Time KVM Integration

Integration of real time KVM with the Compute service further enhances the vCPU scheduling guarantees that CPU pinning provides by reducing the impact of CPU latency resulting from causes such as kernel tasks running on host CPUs. This functionality is crucial to workloads such as network functions virtualization (NFV), where reducing CPU latency is highly important.

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.

VPN-as-a-Service (VPNaaS)

VPN-as-a-Service allows you to create and manage VPN connections in OpenStack.

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 10 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#1188175

This enhancement adds support for virtual device role tagging. This was added because an instance's operating system may need extra information about the virtual devices it is running on. For example, in an instance with multiple virtual network interfaces, the guest operating system needs to distinguish between their intended usage in order to provision them accordingly.

With this update, virtual device role tagging allows users to tag virtual devices when creating an instance. Those tags are then presented to the instance (along with other device metadata) using the metadata API, and through the config drive (if enabled). For more information, see the chapter `Use Tagging for Virtual Device Identification` in the Red Hat OpenStack Platform 10 Networking Guide:

<https://access.redhat.com/documentation/en/red-hat-openstack-platform/>

BZ#1189551

This update adds the `real time` feature, which provides stronger guarantees for worst-case scheduler latency for vCPUs. This update assists tenants that need to run workloads concerned with CPU execution latency, and that require the guarantees offered by a real time KVM guest configuration.

BZ#1198602

This enhancement allows the `admin` user to view a list of the floating IPs allocated to instances, using the admin console. This list spans all projects in the deployment.

Previously, this information was only available from the command-line.

BZ#1233920

This enhancement adds support for virtual device role tagging. This was

added because an instance's operating system may need extra information about the virtual devices it is running on. For example, in an instance with multiple virtual network interfaces, the guest operating system needs to distinguish between their intended usage in order to provision them accordingly.

With this update, virtual device role tagging allows users to tag virtual devices when creating an instance. Those tags are then presented to the instance (along with other device metadata) using the metadata API, and through the config drive (if enabled). For more information, see the chapter `Use Tagging for Virtual Device Identification` in the Red Hat OpenStack Platform 10 Networking Guide:
<https://access.redhat.com/documentation/en/red-hat-openstack-platform/>

BZ#1249836

With the 'openstack baremetal' utility, you can now specify specific images during boot configuration. Specifically, you can now use the '--deploy-kernel' and '--deploy-ramdisk' options to specify a kernel or ramdisk image, respectively.

BZ#1256850

The Telemetry API (ceilometer-api) now uses apache-wsgi instead of eventlet. When upgrading to this release, ceilometer-api will be migrated accordingly.

This change provides greater flexibility for per-deployment performance and scaling adjustments, as well as straightforward use of SSL.

BZ#1262070

You can now use the director to configure Ceph RBD as a Block Storage backup target. This will allow you to deploy an overcloud where volumes are set to back up to a Ceph target. By default, volume backups will be stored in a Ceph pool called 'backups'.

Backup settings are configured in the following environment file (on the undercloud):

```
/usr/share/openstack-tripleo-heat-templates/environments/cinder-backup.yaml
```

BZ#1279554

Using the RBD backend driver (Ceph Storage) for OpenStack Compute (nova) ephemeral disks applies two additional settings to libvirt:

```
hw_disk_discard : unmap
disk_cachemodes : network=writeback
```

This allows reclaiming of unused blocks on the Ceph pool and caching of network writes, which improves the performance for OpenStack Compute

ephemeral disks using the RBD driver.

Also see <http://docs.ceph.com/docs/master/rbd/rbd-openstack/>

BZ#1283336

Previously, in Red Hat Enterprise Linux OpenStack Platform 7, the networks that could be used on each role was fixed. Consequently, it was not possible to have a custom network topology with any network, on any role.

With this update, in Red Hat OpenStack Platform 8 and higher, any network may be assigned to any role.

As a result, custom network topologies are now possible, but the ports for each role will have to be customized. Review the `environments/network-isolation.yaml` file in `openstack-tripleo-heat-templates` to see how to enable ports for each role in a custom environment file or in `network-environment.yaml`.

BZ#1289502

With this release, the customer requires two factor authentication, to support better security for re-seller use case.

BZ#1290251

With this update, a new feature to enable connecting the overcloud to a monitoring infrastructure adds availability monitoring agents (sensu-client) to be deployed on the overcloud nodes.

To enable the monitoring agents deployment, use the environment file `'/usr/share/openstack/tripleo-heat-templates/environments/monitoring-environment.yaml'` and fill in the following parameters in the configuration YAML file:

`MonitoringRabbitHost`: host where the RabbitMQ instance for monitoring purposes is running

`MonitoringRabbitPort`: port on which the RabbitMQ instance for monitoring purposes is running

`MonitoringRabbitUserName`: username to connect to RabbitMQ instance

`MonitoringRabbitPassword`: password to connect to RabbitMQ instance

`MonitoringRabbitVhost`: RabbitMQ vhost used for monitoring purposes

BZ#1309460

You can now use the director to deploy Ceph RadosGW as your object storage gateway. To do so, include `/usr/share/openstack-tripleo-heat-templates/environments/ceph-radosgw.yaml` in your overcloud deployment. When you use this heat template, the default Object Storage service (swift) will not be deployed.

BZ#1314080

With this enhancement, `heat-manage` now supports a `heat-manage`

`reset_stack_status`` subcommand. This was added to manage situations where ``heat-engine`` was unable to contact the database, causing any stacks that were in-progress to remain stuck due to outdated database information. When this occurred, administrators needed a way to reset the status to allow these stacks to be updated again. As a result, administrators can now use the ``heat-manage reset_stack_status`` command to reset a stuck stack.

BZ#1317669

This update includes a release file to identify the overcloud version deployed with OSP director. This gives a clear indication of the installed version and aids debugging. The overcloud-full image includes a new package (`rhosp-release`). Upgrades from older versions also install this RPM. All versions starting with OSP 10 will now have a release file. This only applies to Red Hat OpenStack Platform director-based installations. However, users can manually the install the `rhosp-release` package and achieve the same result.

BZ#1325680

Typically, the installation and configuration of OVS+DPDK in OpenStack is performed manually after overcloud deployment. This can be very challenging for the operator and tedious to do over a large number of Compute nodes. The installation of OVS+DPDK has now been automated in tripleo. Identification of the hardware capabilities for DPDK were previously done manually, and is now automated during introspection. This hardware detection also provides the operator with the data needed for configuring Heat templates. At present, it is not possible to have the co-existence of Compute nodes with DPDK-enabled hardware and without DPDK-enabled hardware.

The ``ironic`` Python Agent discovers the following hardware details and stores it in a swift blob:

- * CPU flags for hugepages support - If `pse` exists then 2MB hugepages are supported If `pdpe1gb` exists then 1GB hugepages are supported
- * CPU flags for IOMMU - If `VT-d/svm` exists, then IOMMU is supported, provided IOMMU support is enabled in BIOS.
- * Compatible nics - compared with the list of NICs whitelisted for DPDK, as listed here <http://dpdk.org/doc/nics>

Nodes without any of the above-mentioned capabilities cannot be used for the Compute role with DPDK.

- * Operator will have a provision to enable DPDK on Compute nodes.
- * The overcloud image for the nodes identified to be Compute-capable and having DPDK NICs, will have the OVS+DPDK package instead of OVS. It will also have packages ``dpdk`` and ``driverctl``.
- * The device names of the DPDK capable NIC's will be obtained from T-H-T. The PCI address of DPDK NIC needs to be identified from the device name. It is required for whitelisting the DPDK NICs during PCI probe.
- * Hugepages needs to be enabled in the Compute nodes with DPDK.
- * CPU isolation needs to be done so that the CPU cores reserved for DPDK Poll Mode Drivers (PMD) are not used by the general kernel balancing, interrupt handling and scheduling algorithms.
- * On each Compute node with a DPDK-enabled NIC, puppet will configure

the DPDK_OPTIONS for whitelisted NICs, CPU mask, and number of memory channels for DPDK PMD. The DPDK_OPTIONS needs to be set in /etc/sysconfig/openvswitch.

`Os-net-config` performs the following steps:

- * Associate the given interfaces with the dpdk drivers (default as vfio-pci driver) by identifying the pci address of the given interface. The driverctl will be used to bind the driver persistently.
- * Understand the ovs_user_bridge and ovs_dpdk_port types and configure the ifcfg scripts accordingly.
- * The "TYPE" ovs_user_bridge will translate to OVS type OVSUserBridge and based on this OVS will configure the datapath type to `netdev`.
- * The "TYPE" ovs_dpdk_port will translate OVS type OVSDPDKPort and based on this OVS adds the port to the bridge with interface type as `dpdk`
- * Understand the ovs_dpdk_bond and configure the ifcfg scripts accordingly.

On each Compute node with a DPDK-enabled NIC, puppet will perform the following steps:

- * Enable OVS+DPDK in /etc/neutron/plugins/ml2/openvswitch_agent.ini
- [OVS] datapath_type=netdev vhostuser_socket_dir=/var/run/openvswitch
- * Configure vhostuser ports in /var/run/openvswitch to be owned by qemu.

On each controller node, puppet will perform the following steps:

- * Add NUMATopologyFilter to scheduler_default_filters in nova.conf.

As a result, the automation of the above-mentioned enhanced platform awareness has been completed, and verified by QA testing.

BZ#1325682

With this update, IP traffic can be managed by DSCP marking rules attached to QoS policies, which are in turn applied to networks and ports.

This was added because different sources of traffic may require different levels of prioritisation at the network level, especially when dealing with real-time information, or critical control data. As a result, the traffic from the specific ports and networks can be marked with DSCP flags. Note that only Open vSwitch is supported in this release.

BZ#1328830

This update adds support for multiple theme configurations. This was added to allow a user to change a theme dynamically, using the front end. Some use-cases include the ability to toggle between a light and dark theme, or the ability to turn on a high contrast theme for accessibility reasons.

As a result, users can now choose a theme at run time.

BZ#1337782

This release now features Composable Roles. TripleO can now be deployed in a composable way, allowing customers to select what services should

run on each node. This, in turn, allows support for more complex use-cases.

BZ#1337783

Generic nodes can now be deployed during the hardware provisioning phase. These nodes are deployed with a generic operating system (namely, Red Hat Enterprise Linux); customers can then deploy additional services directly on these nodes.

BZ#1343130

The package that contains the `ironic-python-agent` image required the `rhosp-director-images` RPM as a dependency. However, you can use the `ironic-python-agent` image for general OpenStack Bare Metal (`ironic`) usage outside of the Red Hat OpenStack Platform director. This update changes the dependencies so that:

- The `rhosp-director-images` RPM requires the `rhosp-director-images-ipa` RPM
- The `rhosp-director-images-ipa` RPM does not require the `rhosp-director-images` RPM

Users now can install the `ironic-python-agent` image separately.

BZ#1346401

It is now possible to confine 'ceph-osd' instances with SELinux policies. In OSP10, new deployments have SELinux configured in 'enforcing' mode on the Ceph Storage nodes.

BZ#1347371

With this enhancement, RabbitMQ introduces the new HA feature of Queue Master distribution. One of the strategies is `min-masters`, which picks the node hosting the minimum number of masters.

This was added because of the possibility that one of the controllers may become unavailable, with Queue Masters then located on available controllers during queue declarations. Once the lost controller becomes available again, masters of newly-declared queues are not placed with priority to the controller with an obviously lower number of queue masters, and consequently the distribution may be unbalanced, with one of the controllers under significantly higher load in the event of multiple fail-overs.

As a result, this enhancement spreads out the queues across controllers after a controller fail-over.

BZ#1353796

With this update, you can now add nodes manually using the UI.

BZ#1359192

■

With this update, the overcloud image includes the Red Hat Cloud Storage 2.0 version installed.

BZ#1366721

The Telemetry service (ceilometer) now uses gnocchi as its default meter dispatcher back end. Gnocchi is more scalable, and is more aligned to the future direction that the Telemetry service is facing.

BZ#1367678

This enhancement adds `NeutronOVSEnabledFirewallDriver`, a new parameter for configuring the Open vSwitch (OVS) firewall driver in Red Hat OpenStack Platform director.

This was added because the neutron OVS agent supports a new mechanism for implementing security groups: the 'openvswitch' firewall.

`NeutronOVSEnabledFirewallDriver` allows users to directly control which implementation is used:

`hybrid` - configures neutron to use the old iptables/hybrid based implementation.

`openvswitch` - enables the new flow-based implementation.

The new firewall driver includes higher performance and reduces the number of interfaces and bridges used to connect guests to the project network. As a result, users can more easily evaluate the new security group implementation.

BZ#1368218

With this update, you can now configure Object Storage service (swift) with additional raw disks by deploying the overcloud with an additional environment file, for example:

```
parameter_defaults:
  ExtraConfig:
    SwiftRawDisks:
      sdb:
        byte_size: 2048
        mnt_base_dir: /src/sdb
      sdc:
        byte_size: 2048
```

As a result, the Object Storage service is not limited by the local node `root` filesystem.

BZ#1371649

This enhancement updates the main script on `sahara-image-element` to only allow the creation of images for supported plugins. For example, you can use the following command to create a CDH 5.7 image using Red Hat Enterprise Linux 7:

```
----
>> ./diskimage-create/diskimage-create.sh -p cloudera -v 5.7
```

Usage: diskimage-create.sh

```
[-p cloudera|mapr|ambari]
[-v 5.5|5.7|2.3|2.4]
[-r 5.1.0]
```

```
----
```

BZ#1381628

As described in <https://bugs.launchpad.net/tripleo/+bug/1630247>, the Sahara service in upstream Newton TripleO is now disabled by default. As part of the upgrade procedure from Red Hat OpenStack Platform 9 to Red Hat OpenStack Platform 10, the Sahara services are enabled/retained by default. If the operator decides they do not want Sahara after the upgrade, they need to include the provided `-e 'major-upgrade-remove-sahara.yaml'` environment file as part of the deployment command for the controller upgrade and converge steps. Note: this environment file must be specified last, especially for the converge step, but it could be done for both steps to avoid confusion. In this case, the Sahara services would not be restarted after the major upgrade.

This approach allows Sahara services to be properly handled during the OSP9 to OSP10 upgrade. As a result, Sahara services are retained as part of the upgrade. In addition, the operator can still explicitly disable Sahara, if necessary.

BZ#1383779

You can now use node-specific hiera to deploy Ceph storage nodes which do not have the same list of block devices. As a result, you can use node-specific hiera entries within the overcloud deployment's Heat templates to deploy non-similar OSD servers.

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#1381227

This update contains the necessary components for testing the use of containers in OpenStack. This feature is available in this release 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#1377763

With Gnocchi 2.2, job dispatch is coordinated between controllers using Redis. As a result, you can expect improved processing of Telemetry measures.

BZ#1385368

To accommodate composable services, NFS mounts used as an Image Service (glance) back end are no longer managed by Pacemaker. As a result, the glance NFS back end parameter interface has changed: The new method is to use an environment file to enable the glance NFS backend. For example:

```
----  
parameter_defaults:  
  GlanceBackend: file  
  GlanceNfsEnabled: true  
  GlanceNfsShare: IP:/some/exported/path  
----
```

Note: the GlanceNfsShare setting will vary depending on your deployment. In addition, mount options can be customized using the ``GlanceNfsOptions`` parameter. If the Glance NFS backend was previously used in Red Hat OpenStack Platform 9, the environment file contents must be updated to match the Red Hat OpenStack Platform 10 format.

3.1.4. Known Issues

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

BZ#1204259

Glance is not configured with `glance.store.http.Store` as a `known_store` in `/etc/glance/glance.conf`. This means the glance client can not create images with the `--copy-from` argument. These commands fail with a "400 Bad Request" error. As a workaround, edit `/etc/glance/glance-api.conf`, add `glance.store.http.Store` to the list in the "stores" configuration option, then restart the `openstack-glance-api` server. This enables successful creation of glance images with the `--copy-from` argument.

BZ#1239130

The director does not provide network validation before or during a deployment. This means a deployment with a bad network configuration can run for two hours with no output and can result in failure. A network validation script is currently in development and will be released in the future.

BZ#1241644

When `openstack-cinder-volume` uses an LVM backend and the Overcloud nodes reboot, the file-backed loopback device is not recreated. As a workaround, manually recreate the loopback device:

```
$ sudo losetup /dev/loop2 /var/lib/cinder/cinder-volumes
```

Then restart `openstack-cinder-volume`. Note that `openstack-cinder-volume` only runs on one node at a time in a high availability cluster of Overcloud Controller nodes. However, the loopback device should exist on all nodes.

BZ#1243306

Ephemeral storage is hard coded as true when using the `NovaEnableRbdBackend` parameter. This means `NovaEnableRbdBackend` instances cannot use cinder backed onto Ceph Storage. As a workaround, add the following to `puppet/hieradata/compute.yaml`:

```
nova::compute::rbd::ephemeral_storage: false
```

This disables ephemeral storage.

BZ#1245826

The `"openstack overcloud update stack"` command returns immediately despite ongoing operations in the background. The command seems to run forever because it's not interactive. In these situations, run the command with the `"-i"` flag. This prompts the user for any manual interaction needs.

BZ#1249210

A timing issue sometimes causes Overcloud neutron services to not automatically start correctly. This means instances are not accessible. As a workaround, you can run the following command on the Controller node cluster:

```
$ sudo pcs resource debug-start neutron-l3-agent
```

Instances will work correctly.

BZ#1266565

Currently, certain setup steps require a SSH connection to the overcloud controllers, and will need to traverse VIPs to reach the Overcloud nodes.

If your environment is using an external load balancer, then these steps are not likely to successfully connect. You can work around this issue by configuring the external load balancer to forward port 22. As a result, the SSH connection to the VIP will succeed.

BZ#1269005

In this release, RHEL OpenStack Platform director only supports a High Availability (HA) overcloud deployment using three controller nodes.

BZ#1274687

■

There is currently a known requirement that can arise when Director connects to the Public API to complete the final configuration post-deployment steps: The Undercloud node must have a route to the Public API, and it must be reachable on the standard OpenStack API ports and port 22 (SSH).

To prepare for this requirement, check that the Undercloud will be able to reach the External network on the controllers, as this network will be used for post-deployment tasks. As a result, the Undercloud can be expected to successfully connect to the Public API after deployment, and perform final configuration tasks. These tasks are required in order for the newly created deployment to be managed using the Admin account.

BZ#1282951

When deploying Red Hat OpenStack Platform director, the bare-metal nodes should be powered off, and the ironic `node-state` and `provisioning-state` must be correct.`

For example, if ironic lists a node as "Available, powered-on", but the server is actually powered off, the node cannot be used for deployment. As a result, you will need to ensure that the node state in ironic matches the actual node state. Use "ironic node-set-power-state <node> [on|off]" and/or "ironic node-set-provisioning-state <node> available" to make the power state in ironic match the real state of the server, and ensure that the nodes are marked `Available`.`

As a result, once the state in ironic is correct, ironic will be able to correctly manage the power state and deploy to the nodes.

BZ#1293379

There is currently a known issue where network configuration changes can cause interface restarts, resulting in an interruption of network connectivity on overcloud nodes.

Consequently, the network interruption can cause outages in the pacemaker controller cluster, leading to nodes being fenced (if fencing is configured). As a result, tripleo-heat-templates is designed to not apply network configuration changes on overcloud updates. By not applying any network configuration changes, the unintended consequence of a cluster outage is avoided.

BZ#1293422

IBM x3550 M5 servers require firmware with minimum versions to work with Red Hat OpenStack Platform.

Consequently, older firmware levels must be upgraded prior to deployment. Affected systems will need to upgrade to the following versions (or newer):

DSA 10.1, IMM2 1.72, UEFI 1.10, Bootcode NA, Broadcom GigE 17.0.4.4a

After upgrading the firmware, deployment should proceed as expected.

BZ#1302081

Address ranges entered for the `AllocationPools` IPv6 networks and IP allocation pools must be input in a valid format according to RFC 5952.`

Consequently, invalid entries will result in an error. As a result, IPv6 addresses should be entered in a valid format: Leading zeros can be omitted or entered in full, and repeating sequences of zeros may be replaced by "::". For example, an IP address of "fd00:0001:0000:0000:00a1:00b2:00c3:0010" may be represented as: "fd00:1::a1:b2:c3:10", but not as: "fd00:01::0b2:0c3:10", because there are an invalid number of leading zeros (01, 0b2, 0c3). The field must be truncated of leading zeros or fully padded.

BZ#1312155

The controller_v6.yaml template contains a parameter for a Management network VLAN. This parameter is not supported in the current version of the director, and can be safely ignored along with any comments referring to the Management network. The Management network references do not need to be copied to any custom templates.

This parameter will be supported in a future version.

BZ#1323024

A puppet manifest bug incorrectly disables LVM partition automounting during the undercloud installation process. As a result, it is possible for undercloud hosts with partitions other than root and swap (activated on kernel command line) to only boot into an emergency shell.

There are several ways to work around this issue. Choose one from the following:

1. Remove the mountpoints manually from /etc/fstab. Doing so will prevent the issue from manifesting in all future cases. Other partitions could also be removed, and the space added to other partitions (like root or swap).
2. Configure the partitions to be activated in /etc/lvm.conf. Doing so will work until the next update/upgrade, when the undercloud installation is re-run.
3. Restrict initial deployment to only root and swap partitions. This will avoid the issue completely.

BZ#1368279

When using Red Hat Ceph as a back end for ephemeral storage, the Compute service does not calculate the amount of available storage correctly. Specifically, Compute simply adds up the amount of available storage without factoring in replication. This results in grossly overstated available storage, which in turn could cause unexpected storage oversubscription.

To determine the correct ephemeral storage capacity, query the Ceph service directly instead.

BZ#1372804

Previously, the Ceph Storage nodes use the local filesystem formatted with `ext4` as the back end for the `ceph-osd` service.

Note: Some `overcloud-full` images for Red Hat OpenStack Platform 9 (Mitaka) were created using `ext4` instead of `xfs`.

With the Jewel release, `ceph-osd` checks the maximum file name length allowed by the back end and refuses to start if the limit is lower than the one configured for Ceph itself. As a workaround, it is possible to verify the filesystem in use for `ceph-osd` by logging on the Ceph Storage nodes and using the following command:

```
# df -l --output=fstype /var/lib/ceph/osd/ceph-$ID
```

Here, \$ID is the OSD ID, for example:

```
# df -l --output=fstype /var/lib/ceph/osd/ceph-0
```

Note: A single Ceph Storage node might host multiple `ceph-osd` instances, in which case there will be multiple subdirectories in `/var/lib/ceph/osd/` for each instance.

If *any* of the OSD instances is backed by an `ext4` filesystem, it is necessary to configure Ceph to use shorter file names, which is possible by deploying/upgrading with an additional environment file, containing the following:

```
parameter_defaults:
  ExtraConfig:
    ceph::profile::params::osd_max_object_name_len: 256
    ceph::profile::params::osd_max_object_namespace_len: 64
```

As a result, you can now verify if each and every `ceph-osd` instance is up and running after an upgrade from Red Hat OpenStack Platform 9 to Red Hat OpenStack Platform 10.

BZ#1383627

Nodes that are imported using "openstack baremetal import --json instackenv.json" should be powered off prior to attempting import. If the nodes are powered on, Ironic will not attempt to add the nodes or attempt introspection.

As a workaround, power off all overcloud nodes prior to running "openstack baremetal import --json instackenv.json".

As a result, if the nodes are powered off, the import should work successfully.

BZ#1383930

If using DHCP HA, the `NeutronDhcpAgentsPerNetwork` value should be set either equal to the number of dhcp-agents, or 3 (whichever is lower), using composable roles. If this is not done, the value will default to

``ControllerCount`` which may not be optimal as there may not be enough `dhcp-agents` running to satisfy spawning that many DHCP servers for each network.

BZ#1385034

When upgrading or deploying a Red Hat OpenStack Platform environment integrated with an external Ceph Storage Cluster from an earlier version (that is, Red Hat Ceph Storage 1.3), you need to enable backwards compatibility. To do so, add an environment file containing the following snippet to your upgrade/deployment:

```
parameter_defaults:
  ExtraConfig:
    ceph::conf::args:
      client/rbd_default_features:
        value: "1"
```

BZ#1391022

Red Hat Enterprise Linux 6 only contains GRUB Legacy, while OpenStack bare metal provisioning (`ironic`) only supports the installation of GRUB2. As a result, deploying a partition image with local boot will fail during the bootloader installation.

As a workaround, if using RHEL 6 for bare metal instances, do not set `boot_option` to `local` in the flavor settings. You can also consider deploying a RHEL 6 whole disk image which already has GRUB Legacy installed.

BZ#1396308

When deploying or upgrading to a Red Hat OpenStack 10 environment that uses Ceph and dedicated blockstorage nodes for LVM, creating instances with attached volumes will no longer work. This is caused by a bug in the way the director configures the Block Storage service during upgrades.

Specifically, the heat templates do not account by default for cases where Ceph and dedicated blockstorage nodes are configured together. As such, the director fails to define some required settings.

Note that LVM is not a suitable Block Storage back end in production, particularly in enterprise environments.

To work around this add an environment file to your upgrade/deployment that contains the following:

```
parameter_defaults:
  BlockStorageExtraConfig:
    tripleo::profile::base::cinder::volume::cinder_enable_iscsi_backend:
      true
    tripleo::profile::base::cinder::volume::cinder_enable_rbd_backend:
      false
```

BZ#1463059

When using Red Hat Ceph Storage as a back end for both Block Storage (cinder) volumes and backups, any attempt to perform an incremental backup will result in a full backup instead, without any warning.

BZ#1321179

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

3.1.5. Deprecated Functionality

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

BZ#1261539

Support for nova-network is deprecated as of Red Hat OpenStack Platform 9 and will be removed in a future release. When creating new environments, it is recommended to use OpenStack Networking (Neutron).

BZ#1404907

In accordance with the upstream project, the LBaaS v1 API has been removed. Red Hat OpenStack Platform 10 supports only the LBaaS v2 API.

3.2. RED HAT OPENSTACK PLATFORM 10 MAINTENANCE RELEASES - MAY 2018 UPDATE

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#1578625

This enhancement backports collectd 5.8 to Red Hat OpenStack Platform 10. collectd 5.8 includes some additional features, such as ovs-events, ovs-stats, and extended libvirt statistics.

* ovs-events: The ovs-events 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. This plugin uses OVS database to get a link state change notification. For more information, see [Plugin ovs_events](#)

* `ovs-stats`: The `ovs-stats` plugin collects statistics of OVS connected interfaces. This plugin uses OVSDDB management protocol (RFC7047) monitor mechanism to get statistics from OVSDDB. For more information, see [Plugin `ovs_stats`](#).

* `Extended libvirt`: The `libvirt` plugin extended to support CMT, MBM, CPU Pinning, utilization, state metrics on the platform. By default, no extra statistics are reported. If enabled, the plugin reports more detailed statistics about the behaviour of virtual machines. For more information, see [Plugin `virt`](#).

* `hugepages`: `collectd` reads the `/sys/devices/system/node/*/hugepages` and `/sys/kernel/mm/hugepages` directories to collect metrics on hugepages. By default, this option is enabled. For more information, see [Plugin `hugepages`](#).

* `rdt`: The `intel_rdt` plugin collects information provided by monitoring features of Intel Resource Director Technology (Intel(R) RDT) like Cache Monitoring Technology (CMT), Memory Bandwidth Monitoring (MBM). These features provide information about utilization of shared resources. For more information, see [Plugin `intel_rdt`](#).

NOTE: To get the latest `collectd` packages, enable the `opstools` repository by running the following command:

```
$ sudo subscription-manager repos --enable=10-optools
```

`collectd` is available as a Technology Preview in this release. For more information on the support scope for features marked as technology previews, see [Technology Preview Features Support Scope](#).

BZ#1258832

With this release, it is now possible to deploy neutron with the OpenDaylight ML2 driver and OpenDaylight L3 DVR service plugin (no OVS Agent or Neutron L3 Agent needed). A pre-defined environment file is provided for OpenDaylight deployments and can be found in ``environments/neutron-.opendaylight-l3.yaml``

Note: The OpenDaylight controller itself is deployed and activated on the first overcloud controller node with default roles. OpenDaylight can also be deployed on a custom role. In addition, with this release there is no support for clustering of the OpenDaylight controller, so only a single instance may be deployed.

BZ#1315651

The High Availability architecture in this release is more simplified, resulting in a less invasive process when services need to be restarted. During scaling operations, only the needed services are restarted. Previously, a scaling operation required the entire cluster to be restarted.

BZ#1337656

The OpenStack Data Processing service now supports version 2.3 of the HDP (Ambari) plug-in.

BZ#1365857

In this release, the Red Hat OpenDaylight is available as a Technology Preview. This version is based on OpenDaylight Boron SR2.

BZ#1365865

The Red Hat OpenDaylight controller does not support clustering in this release, but High Availability is provided for the neutron API service by default.

BZ#1365874

Red Hat OpenDaylight now supports tenant-configurable security groups for IPv4 traffic. In the default setting, each tenant uses a security group that allows communication among instances associated with that group. Consequently, all egress traffic within the security group is allowed, while the ingress traffic from the outside is dropped.

BZ#1415828

This enhancement implements ProcessMonitor in the HaproxyNSDriver class (v2) to use the external_process module, which allows it to monitor and respawn the haproxy processes as needed. The LBaaS agent (v2) will load options related to external_process in order to take a configured action when the Haproxy process dies unexpectedly.

BZ#1415829

This enhancement adds the ability to automatically reschedule load balancers from dead LBaaS agents. Previously, load balancers could be scheduled across multiple LBaaS agents, however if a hypervisor died, the load balancers scheduled to that node would cease operation. With this update, these load balancers are automatically rescheduled to a different agent.

This feature is turned off by default and controlled using ``allow_automatic_lbaas_agent_failover``.

BZ#1469453

Previously, during any stack update operation, there was a unique identifier parameter value called DeployIdentifier that was set to a new timestamp value on every run. This caused puppet to be reapplied across all nodes in the deployment.

This fix adds a new cli arg to "openstack overcloud deploy" called `--skip-deploy-identifier`. The new CLI argument will skip setting this DeployIdentifier value, and puppet will no longer be forced to execute on every stack update.

In some scenarios, Puppet will still execute even if `--skip-deploy-identifier` is passed. Those scenarios include a change to the puppet manifest itself.

Performance of stack update operations, such as scale out, are greatly improved when passing `--skip-deploy-identifier` argument since puppet does not have to run.

BZ#1480338

With this update, the `OS::Nova::ServerGroup` resource now allows the 'soft-affinity' and 'soft-anti-affinity' policies to be used in addition to the 'affinity' and 'anti-affinity' policies.

BZ#1488390

This update adds support for multiple Availability Zones within a single Block Storage (cinder) volume service; this is done by defining the AZ in each driver section.

BZ#1498513

With this update, the ``OS::Neutron::Port`` resource now supports the 'baremetal' and 'direct-physical' (passthrough) `vnic_type`.

BZ#1503896

This release adds support for deploying Dell EMC VMAX Block Storage backend using the Red Hat OpenStack Platform director.

BZ#1508030

This update increases the default value of ``fs.inotify.max_user_instances`` to 1024. This update also allows you to manage the value through a heat template, using ``InotifyIntancesMax``.

BZ#1519867

This update adds support for remote snapshot attachment for backups. This was added because some backends can attach remote snapshots, but have inefficient ``create volume from snapshot`` operations, so the inability to remotely attach a snapshot prevents efficient snapshot or in-use volume backups when scaling out the backup service. As a result, the backup service can efficiently scale out without having to be concerned about whether it is co-located with the volume service on the same node.

BZ#1547323

This feature adds the `rpc_response_timeout` option to the `/etc/cinder/cinder.conf` file.

This adds the ability to configure Cinder's RPC response timeout.

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#1403914

The Dashboard 'Help' button now directs users to the Red Hat OpenStack Platform documentation page (namely, <https://access.redhat.com/documentation/en/red-hat-openstack-platform/>).

BZ#1451714

Problem in detail:

In OSP10 (OvS2.5), following are the issues:

- 1) tuned is configured with wrong set of CPUs. Expected configuration is NeutronDpdkCoreList + NovaVcpuPinSet, but it has been configured as HostCpusList.
- 2) In post-config, the -l of DPDK_OPTIONS is set as 0 and NeutronDpdkCoreList is configured as pmd-cpu-mask

What needs to be corrected after update, manually?

- 1) Add the list of cpus to be isolated, which is NeutronDpdkCoreList + NovaVcpuPinSet to the tuned conf file.

```
TUNED_CORES="<list of CPUs"
sed -i 's/^isolated_cores=.*\/isolated_cores=$TUNED_CORES/'
$tuned_conf_path
tuned-adm profile cpu-partitioning
```

- 2) lcore mask after the update will be set to 0. Get the cpu mask with get_mask code from the first-boot script [1].

```
LCORE_MASK="<mask value output of get_mask"
ovs-vsctl --no-wait set Open_vSwitch . other-config:dppk-lcore-
mask=$LCORE_MASK
```

BZ#1454624

Workaround: Before you upgrade or update OpenStack, delete the guest that attached to the PF. Then you can proceed to update or upgrade and it will pass.

3.2.3. Known Issues

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

BZ#1295374

It is currently not possible to establish the Red Hat OpenStack Platform Director 10 with VxLAN over VLAN tunneling as the VLAN port is not compatible with the DPDK port.

As a workaround, after deploying the Red Hat OpenStack Platform Director with VxLAN, run the following:

```
# ifup br-link
# systemctl restart neutron-openvswitch-agent

* Add the local IP addr to br-link bridge
# ip addr add <local_IP/PREFIX> dev br-link

* Tag br-link port with the VLAN used as tenant network VLAN ID.
# ovs-vsctl set port br-link tag=<VLAN-ID>
```

BZ#1366356

When using userspace datapath (DPDK), some non-PMD threads run on the same CPU that runs PMD (configured by ``pmd-cpu-mask``). This causes the PMD to be preempted which causes latency spikes, drops, etc.

With this update, a fix is implemented within the `post-install.yaml` files available at: <https://access.redhat.com/documentation/en/red-hat-openstack-platform/10/single/network-functions-virtualization-configuration-guide/#ap-ovsdpdk-post-install>.

BZ#1390065

When using OVS-DPDK, all bridges should be of type `ovs_user_bridge` on the Compute node. Red Hat OpenStack Platform director does not support mixing `ovs_bridge` and `ovs_user_bridges` as it kills the OVS-DPDK performances.

BZ#1394402

In order to reduce any interruptions to the allocated CPUs while running either Open vSwitch, virtual machine CPUs or the VNF threads within the virtual machines as much as possible, CPUs should be isolated. However, CPUAffinity cannot prevent all kernel threads from running on these CPUs. To prevent most of the kernel threads, you must use the boot option `'isolcpus=<cpulist>'`. This uses the same CPU list as `'nohz_full'` and `'rcu_nocbs'`. The `'isolcpus'` is engaged right at the kernel boot, and can thus prevent many kernel threads from being scheduled on the CPUs. This could be run on both the hypervisor and guest server.

```
#!/bin/bash
isol_cpus=`awk '{ for (i = 1; i <= NF; i++) if ($i ~ /nohz/) print $i
};'
/proc/cmdline | cut -d"=" -f2`

if [ ! -z "$isol_cpus" ]; then
    grubby --update-kernel=grubby --default-kernel --
args=isolcpus=$isol_cpus
```

```
fi
```

2) The following snippet re-pins the emulator thread action and is not recommended unless you experience specific performance problems.

```
#!/bin/bash
cpu_list=`grep -e "^CPUAffinity=.*" /etc/systemd/system.conf | sed -e
's/CPUAffinity=//' -e 's/ /,/ '`
if [ ! -z "$cpu_list" ]; then
    virsh_list=`virsh list| sed -e '1,2d' -e 's/\s\+/ /g' | awk -F" "
'{print $2}'`
    if [ ! -z "$virsh_list" ]; then
        for vm in $virsh_list; do virsh emulatorpin $vm --cpulist
$cpu_list; done
    fi
fi
```

BZ#1394537

After a `tuned` profile is activated, `tuned` service must start before the `openvswitch` service does, in order to set the cores allocated to the PMD correctly.

As a workaround, you can change the `tuned` service by running the following script:

```
#!/bin/bash

tuned_service=/usr/lib/systemd/system/tuned.service

grep -q "network.target" $tuned_service
if [ "$?" -eq 0 ]; then
    sed -i '/After=.*s/network.target//g' $tuned_service
fi

grep -q "Before=.*network.target" $tuned_service
if [ ! "$?" -eq 0 ]; then
    grep -q "Before=.*" $tuned_service
    if [ "$?" -eq 0 ]; then
        sed -i 's/^\(Before=.*\)\/\1 network.target openvswitch.service/g'
$tuned_service
    else
        sed -i '/After/i Before=network.target openvswitch.service'
$tuned_service
    fi
fi

systemctl daemon-reload
systemctl restart openvswitch
exit 0
```

BZ#1398323

The 'stack delete' command does not delete the mistral environment and

swift container corresponding to the deleted stack.

Use "openstack overcloud plan delete" after deleting a stack.

BZ#1404749

During an upgrade from Red Hat OpenStack Platform (RHOSP) version 9 to version 10, credentials from RHOSP 9 are carried over until convergence, when the full upgrade is completed. This causes alarm evaluation to fail.

Manually update options in '[service_credentials]' section:

1. Set auth_type to password:

```
auth_type=password
```

2. os_* options are no longer valid. Remove os_* prefix from the following options:

```
os_username      - replace with username
os_tenant_name   - replace with project_name
os_password      - replace with password
os_auth_url      - replace with auth_url
os_region_name   - replace with region_name
```

3. Remove 'v2.0' version from auth_url

```
auth_url=http://[fd00:fd00:fd00:2000::10]:5000/
```

4. Restart the service: `systemctl restart openstack-aodh-evaluator.service`

Aodh alarms will now be evaluated correctly.

BZ#1416070

Currently, the Red Hat OpenStack Platform director 10 with SR-IOV overcloud deployment fails when using the NIC IDs (for example, nic1, nic2, nic3 and so on) in the compute.yaml file.

As a workaround, you need to use NIC names (for example, ens1f0, ens1f1, ens2f0, and so on) instead of the NIC IDs to ensure the overcloud deployment completes successfully.

BZ#1416421

While creating the DPDK bond, ``if-up`` of the bond interface will activate the member interfaces by itself. Individual members should not be able to call for ``if-up``. As a result, the deployment fails with bonding in the OVS-DPDK use case.

As a workaround, you need to comment out the interfaces in the ``impl_ifcfg.py`` file as follows:

```
# if base_opt.primary_interface_name:
#     primary_name = base_opt.primary_interface_name
#     self.bond_primary_ifaces[base_opt.name] = primary_name
```

BZ#1481821

The default value of ``pg_num`` and ``pgp_num`` has been set to 128 instead of 32.

Consequently, the existing Ceph pools will be updated so that their ``pg_num`` and ``pgp_num`` changes to 128 and the data will be rebalanced on the OSDs. Customized values previously set in custom Heat environment files will be preserved. To keep ``pg_num`` and ``pgp_num`` set to their previous default values, add an extra environment file to the update or upgrade command. The command should have the following contents:

```
parameter_defaults:
  ExtraConfig:
    ceph::profile::params::osd_pool_default_pg_num: 32
    ceph::profile::params::osd_pool_default_pgp_num: 32
```

BZ#1488517

RHEL overcloud images contain tuned version 2.8.

In OVS-DPDK and SR-IOV deployments, tuned install and activation is done through the first-boot mechanism.

This install and activation fails, as described in https://bugzilla.redhat.com/show_bug.cgi?id=1488369#c1

You need to reboot the compute node to enforce the tuned profile.

BZ#1489070

The new iptables version that ships with RHEL 7.4 includes a new `--wait` parameter. This parameter allows iptables commands issued in parallel to wait until a lock is released by the prior command. For OpenStack, the neutron service provides the iptables locking but only on the routers level.

As such, when processing routers (for example, during a fullsync after the l3 agent is started), some iptables commands issued by neutron may fail because they are experiencing this lock and require the `--wait` parameter that is not available in neutron yet. Any routers affected by this will cause malfunctions of some floating IPs, or some instances may not access the metadata API during cloud-init.

We recommend that you do not upgrade to RHEL 7.4 until neutron is released with a fix that adopts the new iptables `--wait` parameter.

BZ#1549694

Deployments with OVS-DPDK experience a performance degradation, with the following package versions:

```
OVS: openvswitch-2.6.1-16.git20161206.el7ost.x86_64
kernel: 3.10.0-693.17.1.el7.x86_64.
```

3.2.4. Deprecated Functionality

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

BZ#1402497

Certain CLI arguments are considered deprecated and should not be used. The update will allow you to use the CLI args, but there is still a need to specify at the least an environment file to set the ``sat_repo``. You can use an ``env`` file to work around the issue, before running the overcloud command:

1. `cp -r /usr/share/openstack-tripleo-heat-templates/extraconfig/pre_deploy/rhel-registration .`
2. Edit the `rhel-registration/environment-rhel-registration.yaml` and set the `rhel_reg_org`, `rhel_reg_activation_key`, `rhel_reg_method`, `rhel_reg_sat_repo` and `rhel_reg_sat_url` according to your environment.
3. Run the deployment command with `-e rhel-registration/rhel-registration-resource-registry.yaml -e rhel-registration/environment-rhel-registration.yaml`

This workaround has been checked for both Red Hat Satellite 5 and 6, with repos present on the overcloud nodes upon successful deployment.

3.3. RED HAT OPENSTACK PLATFORM 10 MAINTENANCE RELEASE 26 JUNE 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#1365865

The Red Hat OpenDaylight controller does not support clustering in this release, but High Availability is provided for the neutron API service by default.

BZ#1568355

The `dppkvhostuserclient` mode support has been backported. This feature allows OVS to connect to the vhost socket as a client, which allows for reconnecting to the socket without restarting the VM (if OVS crashes or restarts).

NOTE:

- * All VMs should be migrated to `dppkvhostuserclient` mode

- * Live Migration does not work for the existing VM, use either snapshot and create or cold-migration to move to dpdkvhostuserclient mode.
- * Add/Modify the parameter `NeutronVhostuserSocketDir` to `"/var/lib/vhost_sockets"`.
- * Also for a new installation, remove the `"set_ovs_config"` section in the sample `first-boot` script[1].
- * Add the additional environment file `environments/ovs-dpdk-permissions.yaml` for OVS-DPDK deployments (for new installations and minor updates).
- * All these validations are done with OVS version 2.9.

[1] https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/10/html-single/network_functions_virtualization_configuration_guide/index#ap-ovsdpdk-first-boot

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#1587957

During the start of the deployment with `ovs2.9`, the OVS-DPDK is enabled using `first-boot` before applying the kernel arguments and rebooting. This will cause `vswitchd` service to fail initially, but later when the kernel arguments are applied and the node has been rebooted, `vswitchd` will run as expected with DPDK enabled. The failure messages at the initial stage before the reboot can be ignored.

3.3.3. Known Issues

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

BZ#1394402

In order to reduce any interruptions to the allocated CPUs while running either Open vSwitch, virtual machine CPUs or the VNF threads within the virtual machines as much as possible, CPUs should be isolated. However, `CPUAffinity` cannot prevent all kernel threads from running on these CPUs. To prevent most of the kernel threads, you must use the boot option `'isolcpus=<cpulist>'`. This uses the same CPU list as `'nohz_full'` and `'rcu_nocbs'`. The `'isolcpus'` is engaged right at the kernel boot, and can thus prevent many kernel threads from being scheduled on the CPUs. This could be run on both the hypervisor and guest server.

```
#!/bin/bash
```

```
isol_cpus=`awk '{ for (i = 1; i <= NF; i++) if ($i ~ /nohz/) print $i
};'
/proc/cmdline | cut -d"=" -f2`

if [ ! -z "$isol_cpus" ]; then
    grubby --update-kernel=grubby --default-kernel --
args=isolcpus=$isol_cpus
fi
```

2) The following snippet re-pins the emulator thread action and is not recommended unless you experience specific performance problems.

```
#!/bin/bash
cpu_list=`grep -e "^CPUAffinity=.*" /etc/systemd/system.conf | sed -e
's/CPUAffinity=//' -e 's/ /,/ '`
if [ ! -z "$cpu_list" ]; then
    virsh_list=`virsh list | sed -e '1,2d' -e 's/\s\+ /g' | awk -F" "
'{print $2}'`
    if [ ! -z "$virsh_list" ]; then
        for vm in $virsh_list; do virsh emulatorpin $vm --cpulist
$cpu_list; done
    fi
fi
```

BZ#1394537

After a `tuned` profile is activated, `tuned` service must start before the `openvswitch` service does, in order to set the cores allocated to the PMD correctly.

As a workaround, you can change the `tuned` service by running the following script:

```
#!/bin/bash

tuned_service=/usr/lib/systemd/system/tuned.service

grep -q "network.target" $tuned_service
if [ "$?" -eq 0 ]; then
    sed -i '/After=.*s/network.target//g' $tuned_service
fi

grep -q "Before=.*network.target" $tuned_service
if [ ! "$?" -eq 0 ]; then
    grep -q "Before=.*" $tuned_service
    if [ "$?" -eq 0 ]; then
        sed -i 's/^\(Before=.*\)\/\1 network.target openvswitch.service/g'
$tuned_service
    else
        sed -i '/After/i Before=network.target openvswitch.service'
$tuned_service
    fi
fi
```

```
systemctl daemon-reload
systemctl restart openvswitch
exit 0
```

BZ#1489070

The new iptables version that ships with RHEL 7.4 includes a new `--wait` parameter. This parameter allows iptables commands issued in parallel to wait until a lock is released by the prior command. For OpenStack, the neutron service provides the iptables locking but only on the routers level.

As such, when processing routers (for example, during a fullsync after the l3 agent is started), some iptables commands issued by neutron may fail because they are experiencing this lock and require the `--wait` parameter that is not available in neutron yet. Any routers affected by this will cause malfunctions of some floating IPs, or some instances may not access the metadata API during cloud-init.

We recommend that you do not upgrade to RHEL 7.4 until neutron is released with a fix that adopts the new iptables `--wait` parameter.

3.3.4. Deprecated Functionality

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

BZ#1402497

Certain CLI arguments are considered deprecated and should not be used. The update will allow you to use the CLI args, but there is still a need to specify at the least an environment file to set the ``sat_repo``. You can use an ``env`` file to work around the issue, before running the overcloud command:

1. `cp -r /usr/share/openstack-tripleo-heat-templates/extraconfig/pre_deploy/rhel-registration .`
2. Edit the `rhel-registration/environment-rhel-registration.yaml` and set the `rhel_reg_org`, `rhel_reg_activation_key`, `rhel_reg_method`, `rhel_reg_sat_repo` and `rhel_reg_sat_url` according to your environment.
3. Run the deployment command with `-e rhel-registration/rhel-registration-resource-registry.yaml -e rhel-registration/environment-rhel-registration.yaml`

This workaround has been checked for both Red Hat Satellite 5 and 6, with repos present on the overcloud nodes upon successful deployment.

3.4. RED HAT OPENSTACK PLATFORM 10 MAINTENANCE RELEASE 17 SEPTEMBER 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#1559116

With this enhancement, the `OS::Aodh::EventAlarm` Heat resource type is now included in RHEL-OSP 10. This enhancement provides a Heat interface to allow users to define alarms that can be evaluated based on events that other OpenStack services emit. For example; service update, create, or delete events.

BZ#1565295

With this update, bare metal node introspection reports both the `/dev/XXX` block device name and the `/dev/disk/by-path/XXX` name. Unlike the `/dev/XXX` name, the `/dev/disk/by-path/XXX` name does not change at system reboot and may be the same across similarly configured hardware. This update improves reliability of deployments by using `/dev/disk/by-path/XXX` information in the cloud configuration.

BZ#1570949

The hypervisor kernel thread can pre-empt Virtual CPUs (vCPUs) even with strong partitioning enabled (`isolcpus`, `tuned`). These pre-emptions are not frequent, but with 256 descriptors per virtio queue, even one single pre-emption of the vCPU can cause packet drop in network function virtualization (NFC) virtual machines that have a packet rate of 1 Mpps (1 million packets per second) or higher.

With this update, there are two new tunable options for configuring the RX queue size and TX queue size of virtio NICs and reducing packet drop:

- `'rx_queue_size'`
- `'tx_queue_size'`

BZ#1571756

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 that has been 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#1579699

With this enhancement, the Nova 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.

In this update, the restriction of using only the PCID flag is extended to expose multiple CPU feature flags.

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

BZ#1599368

With this update, parallelization of the selinux permission change enables faster upgrade of Ceph OSD.

BZ#1599975

Previously, when an OpenStack service logs at DEBUG level, Oslo Messaging logs the message "Timed out waiting for RPC response" unnecessarily.

With this fix, Oslo Messaging no longer logs this message in instances when the timeout is recoverable.

BZ#1601708

With this update, the hugetlbfs gid value correlates to the kolla fixed gid value to allow easy migration to Red Hat OpenStack Platform 13, where libvirt runs in a kolla container.

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#1589031

With this update, the neutron OVS agent has a new configuration option `bridge_mac_table_size`. This 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) might be overridden by Open vSwitch.

BZ#1608087

When using the linuxbridge ml2 driver, non-privileged tenants can create and attach ports without specifying an IP address, bypassing IP address validation. A potential Denial of Service could occur if an IP address, conflicting with existing guests or routers, is then assigned from outside of the allowed allocation pool.

3.5. RED HAT OPENSTACK PLATFORM 10 MAINTENANCE RELEASE 16 JANUARY 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#1608521

This update introduces vrouter multi-queue support required by a Juniper plug-in. One of Juniper's plug-ins relies on nova to create interfaces with the correct mode - multi-queue or single-queue. vrouter VIFs (OpenContrail) now support multi-queue mode, which allows network performance to scale across multiple vCPUs. To use this feature, create an instance with more than one vCPU from an image that has the `hw_vif_multiqueue_enabled` property set to `true`.

BZ#1653848

This update adds maintenance release version information to the rhosp-release file. For instance:

```
....
$ cat ./etc/rhosp-release
Red Hat OpenStack Platform release 10.0.10 (Newton)
....
```

3.5.2. 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#1613242

This update introduces a new config option, `cpu_pinning_migration_quick_fail`.

The option only applies to instances with pinned CPUs. It controls the failing of live migrations in the scheduler if the required CPUs aren't available on the destination host.

When an instance with CPU pinning is live migrated, the upstream behavior is to keep the CPU mapping on the destination identical to the source. This can result in multiple instances pinned to the same host CPUs. OSP contains a downstream workaround (absent from upstream OpenStack) that prevents live migration if the required CPUs aren't available on the destination host.

The workaround's implementation places the same destination CPU availability restrictions on all other operations that involve scheduling an instance on a new host. These are cold migration, evacuation, resize and unshelve. For these operations, the instance CPU pinning is recalculated to fit the new host, making the restrictions unnecessary. For example, if the exact CPUs needed by an instance are not available on any compute host, a cold migration would fail. Without the workaround, a host would be found that can accept the instance by recalculating its CPU pinning.

You can disable this workaround by setting `cpu_pinning_migration_quick_fail` to `False`. With the quick-fail workaround disabled, live migration with CPU pinning reverts to the upstream behavior, but the restrictions are lifted from all other move operations, allowing them to work correctly.

CHAPTER 4. TECHNICAL NOTES

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

4.1. RHEA-2016:2948 — RED HAT OPENSTACK PLATFORM 10 ENHANCEMENT UPDATE

The bugs contained in this section are addressed by advisory RHEA-2016:2948. Further information about this advisory is available at <https://access.redhat.com/errata/RHEA-2016:2948.html>.

instack-undercloud

BZ#1266509

Previously, `instack-undercloud` did not verify that a subnet mask was provided for the `local_ip` parameter, and incorrectly used a /32 mask. Consequently, networking would not work correctly on the undercloud in this case (for example, introspection would not work). With this update, `instack-undercloud` now validates that a correct subnet mask has been provided.

BZ#1289614

Prior to this update, there was no automated process for periodically purging expired tokens from the Identity Service (keystone) database. Consequently, the keystone database could potentially continue to grow, resulting in a large database size and the possible consumption of all available disk space. With this update, a crontab entry was added to periodically query and delete expired tokens in the keystone database, running once per day. As a result, the keystone database will no longer face unlimited growth due to expired tokens.

BZ#1320318

Previously, the `pxe_ilo` Bare Metal Service (`ironic`) driver would automatically switch to UEFI boot when it detected UEFI-capable hardware, even though the environment might not support UEFI. Consequently, the deployment process failed with `pxe_ilo` drivers when an environment did not support UEFI. With this update, the `pxe_ilo` driver defaults to BIOS boot mode, and a deployment using `pxe_ilo` now works out of the box, regardless of whether UEFI is configured properly.

BZ#1323024

A puppet manifest bug incorrectly disables LVM partition automounting during the undercloud installation process. As a result, it is possible for undercloud hosts with partitions other than root and swap (activated on kernel command line) to only boot into an emergency shell.

There are several ways to work around this issue. Choose one from the following:

1. Remove the mountpoints manually from `/etc/fstab`. Doing so will prevent the issue from manifesting in all future cases. Other partitions could also be removed, and the space added to other partitions (like `root` or `swap`).
2. Configure the partitions to be activated in `/etc/lvm.conf`. Doing so will work until the next update/upgrade, when the undercloud installation is re-run.
3. Restrict initial deployment to only `root` and `swap` partitions. This will avoid the issue completely.

BZ#1324842

Previously, the director auto-generated a value for `'readonly_user_name'` (in `/etc/ceilometer/ceilometer.conf`) that exceeded the 32-characters. This resulted in `ValueSizeConstraint` errors during upgrades. With this release, the director now sets `'readonly_user_name'` to `'ro_snmp_user'` by default, which ensures compliance with the character limit.

BZ#1355818

Previously, the swift proxy pipeline was misconfigured, with the consequence that swift memory usage continued to grow until it was killed. With this fix, proxy-logging has been configured earlier in the swift proxy pipeline. As a result, swift memory usage will not grow continuously.

mariadb-galera

BZ#1375184

Because Red Hat Enterprise Linux 7.3 changed the return format of the `"systemctl is-enabled"` command as consumed by shell scripts, the `mariadb-galera` RPM package, upon installation, erroneously detected that the MariaDB service was enabled when it was not. As a result, the Red Hat OpenStack Platform installer, which then tried to run `mariadb-galera` using `Pacemaker` and not `systemd`, failed to start Galera. With this update, `mariadb-galera`'s RPM installation scripts now use a different `systemctl` command, correctly detecting the default MariaDB as disabled, and the installer can succeed.

BZ#1373598

Previously, both the `'mariadb-server'` and `'mariadb-galera-server'` packages shipped the client-facing libraries: `'dialog.so'` and `'mysql_clear_password.so'`. As a result, the `'mariadb-galera-server'` package would fail to install because of package conflicts.

With this update, the `'dialog.so'` and `'mysql_clear_password.so'` libraries have been moved from `'mariadb-galera-server'` to `'mariadb-libs'`. As a result, the `'mariadb-galera-server'` package installs successfully.

openstack-gnocchi

BZ#1377763

With Gnocchi 2.2, job dispatch is coordinated between controllers using Redis. As a result, you can expect improved processing of Telemetry measures.

openstack-heat

BZ#1349120

Prior to this update, Heat would occasionally consider a ``FloatingIP`` resource deleted while the deletion was in fact still in progress. Consequently, resources that the ``FloatingIP`` depended on would sometimes fail to be deleted because the ``FloatingIP`` still existed. With this update, Heat now checks that the ``FloatingIP`` can no longer be found before considering the resource deleted, and stack deletes should proceed normally.

BZ#1375930

Previously, the ``str_replace`` intrinsic function worked by calling the Python ``str.replace()`` method for each string to be replaced. Consequently, if the replacement text for one replacement contained another of the strings to be replaced, the replacement text itself could be replaced. The result was non-deterministic, since the replacement order was not guaranteed. Therefore users had to be careful to use techniques, such as guard characters, to ensure that there was no misinterpretation.

With this update, replacements are now performed in a single pass, so only the original text is subject to replacement. As a result, the output of ``str_replace`` is now deterministic, and consistent with user expectations even without the use of guard characters. When keys overlap in the input, longer matches are preferred. Lexicographically smaller strings will be replaced first if there is still ambiguity.

BZ#1314080

With this enhancement, ``heat-manage`` now supports a ``heat-manage reset_stack_status`` subcommand. This was added to manage situations where ``heat-engine`` was unable to contact the database, causing any stacks that were in-progress to remain stuck due to outdated database information. When this occurred, administrators needed a way to reset the status to allow these stacks to be updated again.

As a result, administrators can now use the ``heat-manage``

`reset_stack_status`` command to reset a stuck stack.

openstack-ironic

BZ#1347475

This update adds a socat-based serial console for IPMITool drivers. This was added because users may want to access a bare metal node's serial console in the same way that they access a virtual node's console. As a result, the new driver ``pxe_ipmitool_socat`` was added, with support for the serial console using the ``socat`` utility.

BZ#1310883

The Bare Metal provisioning service now wipes a disk's metadata before partitioning and writing an image into it. This ensures that the new image boots normally. In previous releases, the Bare Metal provisioning service didn't remove old metadata before starting work on a device, which made it possible for a deployment to fail.

BZ#1319841

The openstack-ironic-conductor service now checks whether all drivers specified in the 'enabled_drivers' option are unique. The service then removes duplicated entries and logs a warning. In previous releases, duplicate entries in the 'enabled_drivers' option simply caused the openstack-ironic-conductor service to fail, thereby preventing the Bare Metal provisioning service from loading any drivers.

BZ#1344004

Previously, 'ironic-conductor' did not correctly pass the authentication token to the 'python-neutronclient'. As a result, automatic node cleaning failed with a tear down error.

With this update, OpenStack Baremetal Provisioning (ironic) was migrated to use the 'keystoneauth' sessions rather than directly constructing Identity service client objects. As a result, nodes can now be successfully torn down after cleaning.

BZ#1385114

To determine which node is being deployed, the deploy ramdisk (IPA) provides the Bare Metal provisioning service with a list of MAC addresses as unique identifiers for that node. In previous releases, the Bare Metal provisioning service only expected normal MAC address formats; namely, 6 octets. The GUID of Infiniband NICs, however, have 20 octets. As such, whenever an Infiniband NIC was present on the node, the deployment would fail since the Bare Metal provisioning API could not validate the MAC address correctly.

With this release, the Bare Metal provisioning service now ignores MAC addresses that don't conform with the normal MAC address format of 6 octets.

BZ#1387322

This release removes a redundant 'dhcp' command from the iPXE templates for deployment and introspection. In some cases, this redundant command caused an incorrect interface to receive an IP address.

openstack-ironic-inspector**BZ#1323735**

Previously, the modification dates were not being set on the IPA RAM disk logs when creating a tarfile. As a result, the introspection logs appeared to have the modification date of 1970-01-01, causing GNU tar to issue a warning when extracting the files.

With this update, the modification dates are set correctly when creating a tarfile. The timestamps are now correct and GNU tar no longer issues the warning.

openstack-ironic-python-agent**BZ#1393008**

This release features more thorough error checking and handling around LLDP discovery. This enhancement prevents malformed packages from failing LLDP discovery; in addition, failed LLDP discovery no longer fails the whole introspection process.

openstack-manila**BZ#1380482**

Prior to this update, the Manila Ceph FS driver did not check if it could connect to the Ceph server. Consequently, if the connection to the Ceph server did not work, `manila-share` service kept crashing or respawning without any timeout. With this update, there is now a check to confirm that the Ceph connection works when initializing the Manila Ceph FS driver. As a result, the Ceph driver checks the Ceph connection on driver init, and if it fails the driver is not initialized and no further steps are performed.

openstack-neutron**BZ#1381620**

Previously, the maximum number of client connections (i.e greenlets spawned at a time) opened at any time by the WSGI server was set to 100 with 'wsgi_default_pool_size'. While this setting was adequate for the OpenStack Networking API server, the state change server created heavy CPU loads on the L3 agent, which caused the agent to crash.

With this release, you can now use the new 'ha_keepalived_state_change_server_threads' setting to configure the number of threads in the state change server. Client connections are no longer limited by 'wsgi_default_pool_size', thereby avoiding an L3 agent crash when many state change server threads are spawned.

BZ#1382717

Previously, the 'vport_gre' kernel module had a dependency on the 'ip_gre' kernel module in Red Hat Enterprise Linux 7.3. The 'ip_gre' module created two new interfaces: 'gre0' and 'gretap0'. These interfaces are created in each namespace and cannot be removed. As a result, when 'neutron-netns'cleanup' purged all the interfaces during the namespace cleanup, the 'gre0' and 'gretap0' were not removed. This prevented the network namespace from being deleted due to some interfaces still being present.

With this update, the 'gre0' and 'gretap0' interfaces are added to the whitelist of interfaces and are ignored when checking whether the namespace contains any interface. As a result, the network namespace is deleted even when it contains the 'gre0' and 'gretap0' interfaces.

BZ#1384334

This release adds a HTTPProxyToWSGI middleware in front of the OpenStack Networking API to set up a request URL correctly in case a proxy (eg. HAProxy) is used between the client and server. This ensures that when a client uses SSL, the server recognizes this and responds using the correct protocol. Previously, using a proxy made it possible for the server to respond with HTTP (instead of HTTPS) even when a client used SSL.

BZ#1387546

Previously, it was possible for the OpenStack networking OVS agent to compare non-translated string to translated, UTF-16 strings when a subprocess didn't run properly. On non-English locales, this could result in an exception, thereby preventing instances from booting.

To address this, failure checks were updated to depend on the actual return value of failed subprocesses instead of strings. This ensures that subprocess failures are handled properly under non-English locales.

BZ#1325682

With this update, IP traffic can be managed by DSCP marking rules attached to QoS policies, which are in turn applied to networks and

ports.

This was added because different sources of traffic may require different levels of prioritisation at the network level, especially when dealing with real-time information, or critical control data. As a result, the traffic from the specific ports and networks can be marked with DSCP flags. Note that only Open vSwitch is supported in this release.

openstack-nova

BZ#1188175

This enhancement adds support for virtual device role tagging. This was added because an instance's operating system may need extra information about the virtual devices it is running on. For example, in an instance with multiple virtual network interfaces, the guest operating system needs to distinguish between their intended usage in order to provision them accordingly.

With this update, virtual device role tagging allows users to tag virtual devices when creating an instance. Those tags are then presented to the instance (along with other device metadata) using the metadata API, and through the config drive (if enabled). For more information, see the chapter `Use Tagging for Virtual Device Identification` in the Red Hat OpenStack Platform 10 Networking Guide:

<https://access.redhat.com/documentation/en/red-hat-openstack-platform/>

BZ#1189551

This update adds the `real time` feature, which provides stronger guarantees for worst-case scheduler latency for vCPUs. This update assists tenants that need to run workloads concerned with CPU execution latency, and that require the guarantees offered by a real time KVM guest configuration.

BZ#1233920

This enhancement adds support for virtual device role tagging. This was added because an instance's operating system may need extra information about the virtual devices it is running on. For example, in an instance with multiple virtual network interfaces, the guest operating system needs to distinguish between their intended usage in order to provision them accordingly.

With this update, virtual device role tagging allows users to tag virtual devices when creating an instance. Those tags are then presented to the instance (along with other device metadata) using the metadata API, and through the config drive (if enabled). For more information, see the chapter `Use Tagging for Virtual Device Identification` in the Red Hat OpenStack Platform 10 Networking Guide:

<https://access.redhat.com/documentation/en/red-hat-openstack-platform/>

BZ#1263816

Previously, the nova ironic virt driver wrote an instance UUID in the Bare Metal Provisioning (ironic) node before starting a deployment. If something failed between writing the UUID and starting the deployment, Compute did not remove the instance after it failed to spawn the instance. As a result, the Bare Metal Provisioning (ironic) node would have an instance UUID set and would not be picked for another deployment.

With this update, if spawning an instance fails at any stage of the deployment, the ironic virt driver ensures that the instance UUID is cleaned up. As a result, nodes will not have an instance UUID set and will be picked up for a new deployment.

openstack-puppet-modules

BZ#1284058

Previously, Object Storage service deployed using the director used ceilometer middleware that had been deprecated since the Red Hat OpenStack Platform 8 (liberty) release.

With this update, the Object Storage service has been fixed to use the ceilometer middleware from python-ceilometermiddleware which is the supported version for this release.

BZ#1372821

Previously, the Time Series Database-as-a-Service (gnocchi) API workers were configured to be deployed by default with a single process and logical `cpu_core` count for threads, resulting in the gnocchi API running in `httpd` to be deployed with a single process.

As a best practice, gnocchi recommends the number of process and threads to be $1.5 * \text{cpu_count}$. With this update, the worker count is $\max((\$::processorcount + 0)/4, 2)$ and threads to 1. As a result, the gnocchi API workers run with the right number of workers and threads for better performance.

openstack-tripleo-common

BZ#1382174

Previously, the 'DeployIdentifier' was not being updated for package update, resulting in Puppet not being run on the non-controller nodes.

With this update, the 'DeployIdentifier' value is incremented. As a result, Puppet runs and updates packages on the non-controller nodes.

BZ#1323700

Previously, in the OpenStack Director, the 'upgrade-non-controller.sh'

script used by an operator on the Undercloud to upgrade the non-controller nodes as a part of the major upgrade workflow did not report the upgrade status when the '--query' option was used. As a result, the '--query' option did not work as documented by the '-h' helptext.

With this update, the '--query' option now provides the last few lines of the 'yum.log' file from the given node as an indication of the upgrade status. Also, the script now accepts the long and short versions for each of the options ('-q' and '--query'). As a result, the 'upgrade-non-controller.sh' script is now improved to provide at least some indication of the node upgrade status.

BZ#1383627

Nodes that are imported using "openstack baremetal import --json instackenv.json" should be powered off prior to attempting import. If the nodes are powered on, Ironic will not attempt to add the nodes or attempt introspection.

As a workaround, power off all overcloud nodes prior to running "openstack baremetal import --json instackenv.json".

As a result, if the nodes are powered off, the import should work successfully.

openstack-tripleo-heat-templates

BZ#1262064

It is now possible to deploy 'cinder-backup' in the overcloud using a Heat environment file when launching the stack deployment. The environment file which enables 'cinder-backup' is /usr/share/openstack-tripleo-heat-templates/environments/cinder-backup.yaml. The 'cinder-backup' service will initially support the use of Swift or Ceph as backends. The 'cinder-backup' service performs backups of Cinder volumes on backends different than the one where the volumes are stored. The 'cinder-backup' service will be running in the overcloud if included at deployment time.

BZ#1282491

Prior to this update, the RabbitMQ maximum open file descriptors was set to 4096. Consequently, customers with larger deployments could hit this limit and face stability issues. With this update, the maximum open file descriptor limit for RabbitMQ has been increased to 65536. As a result, larger deployments should now be significantly less likely to run into this issue.

BZ#1242593

With this enhancement, the OpenStack Bare Metal provisioning service (ironic) can be deployed in the overcloud to support the provision of bare metal instances. This was added because customers may want to deploy bare metal instances in their overcloud.

As a result, the Red Hat OpenStack Platform director can now optionally deploy the Bare metal service in order to provision bare metal instances in the overcloud.

BZ#1274196

With this update, the iptables firewall on the overcloud controller nodes are enabled to ensure better security. As a result, the necessary ports are opened so that overcloud services will continue to function as before.

BZ#1290251

With this update, a new feature to enable connecting the overcloud to a monitoring infrastructure adds availability monitoring agents (sensu-client) to be deployed on the overcloud nodes.

To enable the monitoring agents deployment, use the environment file `'/usr/share/openstack/tripleo-heat-templates/environments/monitoring-environment.yaml'` and fill in the following parameters in the configuration YAML file:

```
MonitoringRabbitHost: host where the RabbitMQ instance for monitoring
purposes is running
MonitoringRabbitPort: port on which the RabbitMQ instance for
monitoring purposes is running
MonitoringRabbitUserName: username to connect to RabbitMQ instance
MonitoringRabbitPassword: password to connect to RabbitMQ instance
MonitoringRabbitVhost: RabbitMQ vhost used for monitoring purposes
```

BZ#1309460

You can now use the director to deploy Ceph RadosGW as your object storage gateway. To do so, include `/usr/share/openstack-tripleo-heat-templates/environments/ceph-radosgw.yaml` in your overcloud deployment. When you use this heat template, the default Object Storage service (swift) will not be deployed.

BZ#1325680

Typically, the installation and configuration of OVS+DPDK in OpenStack is performed manually after overcloud deployment. This can be very challenging for the operator and tedious to do over a large number of Compute nodes. The installation of OVS+DPDK has now been automated in tripleo. Identification of the hardware capabilities for DPDK were previously done manually, and is now automated during introspection. This hardware detection also provides the operator with the data needed for configuring Heat templates. At present, it is not possible to have the co-existence of Compute nodes with DPDK-enabled hardware and without DPDK-enabled hardware.

The ``ironic`` Python Agent discovers the following hardware details and stores it in a swift blob:

* CPU flags for hugepages support - If `pse` exists then 2MB hugepages are

supported If pdpe1gb exists then 1GB hugepages are supported

- * CPU flags for IOMMU - If VT-d/svm exists, then IOMMU is supported, provided IOMMU support is enabled in BIOS.
- * Compatible nics - compared with the list of NICs whitelisted for DPDK, as listed here <http://dpdk.org/doc/nics>

Nodes without any of the above-mentioned capabilities cannot be used for the Compute role with DPDK.

- * Operator will have a provision to enable DPDK on Compute nodes.
- * The overcloud image for the nodes identified to be Compute-capable and having DPDK NICs, will have the OVS+DPDK package instead of OVS. It will also have packages `dpdk` and `driverctl`.
- * The device names of the DPDK capable NIC's will be obtained from T-H-T. The PCI address of DPDK NIC needs to be identified from the device name. It is required for whitelisting the DPDK NICs during PCI probe.
- * Hugepages needs to be enabled in the Compute nodes with DPDK.
- * CPU isolation needs to be done so that the CPU cores reserved for DPDK Poll Mode Drivers (PMD) are not used by the general kernel balancing, interrupt handling and scheduling algorithms.
- * On each Compute node with a DPDK-enabled NIC, puppet will configure the DPDK_OPTIONS for whitelisted NICs, CPU mask, and number of memory channels for DPDK PMD. The DPDK_OPTIONS needs to be set in `/etc/sysconfig/openvswitch`.

``Os-net-config`` performs the following steps:

- * Associate the given interfaces with the dpdk drivers (default as vfio-pci driver) by identifying the pci address of the given interface. The driverctl will be used to bind the driver persistently.
- * Understand the ovs_user_bridge and ovs_dpdk_port types and configure the ifcfg scripts accordingly.
- * The "TYPE" ovs_user_bridge will translate to OVS type OVSUserBridge and based on this OVS will configure the datapath type to 'netdev'.
- * The "TYPE" ovs_dpdk_port will translate OVS type OVSDPDKPort and based on this OVS adds the port to the bridge with interface type as 'dpdk'
- * Understand the ovs_dpdk_bond and configure the ifcfg scripts accordingly.

On each Compute node with a DPDK-enabled NIC, puppet will perform the following steps:

- * Enable OVS+DPDK in `/etc/neutron/plugins/ml2/openvswitch_agent.ini`
- [OVS] datapath_type=netdev vhostuser_socket_dir=/var/run/openvswitch
- * Configure vhostuser ports in `/var/run/openvswitch` to be owned by qemu.

On each controller node, puppet will perform the following steps:

- * Add NUMATopologyFilter to scheduler_default_filters in nova.conf.

As a result, the automation of the above-mentioned enhanced platform awareness has been completed, and verified by QA testing.

BZ#1337782

This release now features Composable Roles. TripleO can now be deployed in a composable way, allowing customers to select what services should run on each node. This, in turn, allows support for more complex use-cases.

■

BZ#1337783

Generic nodes can now be deployed during the hardware provisioning phase. These nodes are deployed with a generic operating system (namely, Red Hat Enterprise Linux); customers can then deploy additional services directly on these nodes.

BZ#1381628

As described in <https://bugs.launchpad.net/tripleo/+bug/1630247>, the Sahara service in upstream Newton TripleO is now disabled by default. As part of the upgrade procedure from Red Hat OpenStack Platform 9 to Red Hat OpenStack Platform 10, the Sahara services are enabled/retained by default. If the operator decides they do not want Sahara after the upgrade, they need to include the provided `-e 'major-upgrade-remove-sahara.yaml'` environment file as part of the deployment command for the controller upgrade and converge steps. Note: this environment file must be specified last, especially for the converge step, but it could be done for both steps to avoid confusion. In this case, the Sahara services would not be restarted after the major upgrade.

This approach allows Sahara services to be properly handled during the OSP9 to OSP10 upgrade. As a result, Sahara services are retained as part of the upgrade. In addition, the operator can still explicitly disable Sahara, if necessary.

BZ#1389502

This update allows for custom values for the `kernel.pid_max` sysctl key using the `KernelPidMax` Heat parameter with a default of 1048576. On nodes working as Ceph clients there might be a large number of running threads, depending on the number of `ceph-osd` instances. In such cases, the `pid_max` might reach the maximum value and cause I/O errors. The `pid_max` key has a higher default and can be customized via `KernelPidMax` parameter.

BZ#1243483

Previously, polling the Orchestration service for server metadata resulted in REST API calls to Compute, resulting in a constant load on the `nova-api` which worsened as the cloud was scaled up.

With this update, Object Storage service is now polled for server metadata and loading the heat stack no longer makes unnecessary calls to the `nova-api`. As a result, there is a significant reduction in the load on the undercloud as the overcloud scales up.

BZ#1315899

Previously, the director-deployed swift used a deprecated version of ceilometer middleware that had been dropped in Red Hat OpenStack Platform 8. With this update, the swift proxy config uses ceilometer

middleware from python-ceilometermiddleware. As a result, swift proxy now uses a supported version of ceilometer middleware.

BZ#1361285

OpenStack Image Storage (glance) configures with more workers by default, which improves performance. The count is automatically scaled depending on the number of processors.

BZ#1367678

This enhancement adds `NeutronOVSTFirewallDriver`, a new parameter for configuring the Open vSwitch (OVS) firewall driver in Red Hat OpenStack Platform director.

This was added because the neutron OVS agent supports a new mechanism for implementing security groups: the 'openvswitch' firewall.

`NeutronOVSTFirewallDriver` allows users to directly control which implementation is used:

`hybrid` - configures neutron to use the old iptables/hybrid based implementation.

`openvswitch` - enables the new flow-based implementation.

The new firewall driver includes higher performance and reduces the number of interfaces and bridges used to connect guests to the project network. As a result, users can more easily evaluate the new security group implementation.

BZ#1256850

The Telemetry API (ceilometer-api) now uses apache-wsgi instead of eventlet. When upgrading to this release, ceilometer-api will be migrated accordingly.

This change provides greater flexibility for per-deployment performance and scaling adjustments, as well as straightforward use of SSL.

BZ#1303093

With this update, it is possible to disable the Object Storage service (swift) in the overcloud by using an additional environment file when deploying the overcloud. The environment file should contain the following:

```
resource_registry:
  OS::Triple0::Services::SwiftProxy: OS::Heat::None
  OS::Triple0::Services::SwiftStorage: OS::Heat::None
  OS::Triple0::Services::SwiftRingBuilder: OS::Heat::None
```

As a result, the Object Storage service will not be running in the overcloud and there will not be an endpoint for the Object Storage service in the overcloud Identity service.

BZ#1314732

■

Previously, while deploying Red Hat OpenStack Platform 8 using director, the Telemetry service was not configured in Compute, causing some of the OpenStack Integration Test Suite tests to fail.

With this update, the OpenStack Telemetry service is configured in the Compute configuration. As a result, the notification driver is set correctly and the OpenStack Integration Test Suite tests pass.

BZ#1316016

Previously, Telemetry (ceilometer) notifications would fail due to missing messaging configuration in Image Service (glance). Consequently, glance notifications failed to be processed. With this update, the tripleo templates have been amended to add the correct configuration. As a result, glance notifications are now processed correctly.

BZ#1347371

With this enhancement, RabbitMQ introduces the new HA feature of Queue Master distribution. One of the strategies is `min-masters`, which picks the node hosting the minimum number of masters.

This was added because of the possibility that one of the controllers may become unavailable, with Queue Masters then located on available controllers during queue declarations. Once the lost controller becomes available again, masters of newly-declared queues are not placed with priority to the controller with an obviously lower number of queue masters, and consequently the distribution may be unbalanced, with one of the controllers under significantly higher load in the event of multiple fail-overs.

As a result, this enhancement spreads out the queues across controllers after a controller fail-over.

BZ#1351271

The Red Hat OpenStack Platform director creates OpenStack Block Storage (cinder) v3 API endpoint in OpenStack Identity (keystone) to support the newer Cinder API version.

BZ#1364478

This update allow usage of any isolated network on any role. Some scenarios, like a deployment where 'ceph-osd' is collocated with 'nova-compute', assume that nodes have access to multiple isolated networks. Now custom NIC templates can configure any of the isolated network on any role.

BZ#1366721

The Telemetry service (ceilometer) now uses gnocchi as its default meter dispatcher back end. Gnocchi is more scalable, and is more aligned to the future direction that the Telemetry service is facing.

BZ#1368218

With this update, you can now configure Object Storage service (swift) with additional raw disks by deploying the overcloud with an additional environment file, for example:

```
parameter_defaults:
  ExtraConfig:
    SwiftRawDisks:
      sdb:
        byte_size: 2048
        mnt_base_dir: /src/sdb
      sdc:
        byte_size: 2048
```

As a result, the Object Storage service is not limited by the local node ``root`` filesystem.

BZ#1369426

AODH now uses MYSQL as its default database back end. Previously, AODH used MongoDB as its default back end to make the transition from Ceilometer to AODH easier.

BZ#1373853

The Compute role and Object Storage role upgrade scripts for upgrading from the Red Hat OpenStack Platform 9 (mitaka) to Red Hat OpenStack Platform 10 (newton) did not exit on error as expected. As a result, the `'upgrade-non-controller.sh'` script returned code 0 (success) even when the upgrade failed.

With this update, the Compute role and the Object Storage role upgrade scripts now exit on error during the upgrade process and the `'upgrade-non-controller.sh'` returns a non-zero (failure) value if the upgrade fails.

BZ#1379719

With the move to composable services, the hieradata which was used to configure the NTP servers on overcloud nodes was configured incorrectly.

This update uses the correct hieradata so the overcloud nodes get the NTP servers configured.

BZ#1385368

To accommodate composable services, NFS mounts used as an Image Service (glance) back end are no longer managed by Pacemaker. As a result, the glance NFS back end parameter interface has changed: The new method is to use an environment file to enable the glance NFS backend. For example:

```
----
```

```
parameter_defaults:  
GlanceBackend: file  
GlanceNfsEnabled: true  
GlanceNfsShare: IP:/some/exported/path  
-----
```

Note: the GlanceNfsShare setting will vary depending on your deployment. In addition, mount options can be customized using the `GlanceNfsOptions` parameter. If the Glance NFS backend was previously used in Red Hat OpenStack Platform 9, the environment file contents must be updated to match the Red Hat OpenStack Platform 10 format.

BZ#1387390

Previously, the TCP port '16509' was blocked in 'iptables'. As a result, the 'nova' Compute 'libvirt' instances could not be live migrated between Compute nodes.

With this update, TCP port '16509' is configured to be opened in the 'iptables'. As a result, the 'nova' Compute 'libvirt' instances can now be live migrated between Compute nodes.

BZ#1389189

Previously, due to a race condition between Hieradata getting written and Puppet execution on nodes, Puppet on the Overcloud nodes failed occasionally due to the missing Hieradata.

With this update, ordering is introduced, first writing of the Hieradata is completed on all nodes and then Puppet execution takes place. As a result, Puppet no longer fails during execution as all the necessary Hieradata is present.

BZ#1392773

Previously, after upgrading from Red Hat OpenStack Platform 9 (Mitaka) to Red Hat OpenStack Platform 10 (Newton), the 'ceilometer-compute-agent' failed to collect data.

With this update, restarting the 'ceilometer-compute-agent' post upgrade fixes the issue and allows the 'ceilometer-compute-agent' to restart correctly and gather the relevant data.

BZ#1393487

OpenStack Platform director did not update firewall when deploying OpenStack File Share API (manila-api). If you moved the manila-api service off controllers to its own role, the default firewall rules blocked the endpoints. This fix updates the manila-api firewall rules in the overcloud Heat template collection. You can now reach the endpoints even when manila-api is on a role separate from the controller nodes.

BZ#1382579

-

The director set the cloudformation (heat-cfn) endpoint to "RegionOne" instead of "regionOne". This caused the UI to display two regions with different services. This fix sets the endpoint to use "regionOne". The UI now displays all services under the same region.

openstack-tripleo-ui

BZ#1353796

With this update, you can now add nodes manually using the UI.

os-collect-config

BZ#1306140

Prior to this update, HTTP requests to ``os-collect-config`` for configuration did not specify a request timeout. Consequently, polling for data while the undercloud was inaccessible (for example, rebooting undercloud, network connectivity issues) resulted in ``os-collect-config`` stalling, performing no polling or configuration. This often only became apparent when an overcloud stack operation was performed and software configuration operations timed out.

With this update, ``os-collect-config`` HTTP requests now always specify a timeout period.

As a result, polling for data will fail when the undercloud is unavailable, and then resume when it is available again.

os-net-config

BZ#1391031

Prior to this update, improvements in the integration between Open vSwitch and neutron could cause issues with the resumption of connectivity after a restart. Consequently, nodes could become unreachable or have reduced connectivity.

With this update, ``os-net-config`` configures ``fail_mode=standalone`` by default to allow network traffic if no controlling agent has started yet. As a result, the connection issues on reboot have been resolved.

puppet-ceph

BZ#1372804

Previously, the Ceph Storage nodes use the local filesystem formatted with ``ext4`` as the back end for the ``ceph-osd`` service.

Note: Some ``overcloud-full`` images for Red Hat OpenStack Platform 9 (Mitaka) were created using ``ext4`` instead of ``xfs``.

With the Jewel release, ``ceph-osd`` checks the maximum file name length

allowed by the back end and refuses to start if the limit is lower than the one configured for Ceph itself. As a workaround, it is possible to verify the filesystem in use for `ceph-osd` by logging on the Ceph Storage nodes and using the following command:

```
# df -l --output=fstype /var/lib/ceph/osd/ceph-$ID
```

Here, \$ID is the OSD ID, for example:

```
# df -l --output=fstype /var/lib/ceph/osd/ceph-0
```

Note: A single Ceph Storage node might host multiple `ceph-osd` instances, in which case there will be multiple subdirectories in `/var/lib/ceph/osd/` for each instance.

If *any* of the OSD instances is backed by an `ext4` filesystem, it is necessary to configure Ceph to use shorter file names, which is possible by deploying/upgrading with an additional environment file, containing the following:

```
parameter_defaults:
```

```
  ExtraConfig:
```

```
    ceph::profile::params::osd_max_object_name_len: 256
```

```
    ceph::profile::params::osd_max_object_namespace_len: 64
```

As a result, you can now verify if each and every `ceph-osd` instance is up and running after an upgrade from Red Hat OpenStack Platform 9 to Red Hat OpenStack Platform 10.

BZ#1346401

It is now possible to confine 'ceph-osd' instances with SELinux policies. In OSP10, new deployments have SELinux configured in 'enforcing' mode on the Ceph Storage nodes.

BZ#1370439

Reusing Ceph nodes from an previous cluster in a new overcloud caused the new Ceph cluster to fail without any indication during the overcloud deployment process. This was because the old Ceph OSD node disks needed cleaning before reusing them. This fix adds a check to the Ceph OpenStack Puppet module to make sure the disks are clean as per the instructions in the OpenStack Platform documentation [1]. Now the overcloud deployment process properly fails if it detects non-clean OSD disks. The 'openstack stack failures list overcloud' command indicates the disks which have a FSID mismatch.

[1] https://access.redhat.com/documentation/en/red-hat-openstack-platform/10/single/red-hat-ceph-storage-for-the-overcloud/#Formatting_Ceph_Storage_Nodes_Disks_to_GPT

puppet-cinder

BZ#1356683

A race condition existed between loop device configuration and a check for LVM physical volumes on block storage nodes. This caused the major upgrade convergence step to fail due to Puppet being failing to detect existing LVM physical volumes and attempting to recreate the volume. This fix waits for udev events to complete after setting up the loop device. This means that Puppet waits for the loop device configuration to complete before attempting to check for an existing LVM physical volume. Block storage nodes with LVM backends now upgrade successfully.

puppet-heat

BZ#1381561

The OpenStack Platform director exceeded the default memory limits for using OpenStack Orchestration (heat) YAQL expressions. This caused an "Expression consumed too much memory" error during an overcloud deployment and subsequent deployment failure. This fix increases the default memory limits for the director, which results in a error-free overcloud deployment.

puppet-ironic

BZ#1314665

The ironic-inspector server did not have an iPXE version that worked with UEFI bootloaders. Machines with UEFI bootloaders could not chainload the introspection ramdisk. This fix ensures the ipxe.efi ROM is present on the ironic-inspector server and updates the dnsmasq configuration to send it to the UEFI-based machine during introspection. Now the director can inspect both BIOS and UEFI machines.

puppet-tripleo

BZ#1386611

rabbitmqctl failed to function in an IPv6 environment due to a missing parameter. This fix modifies the RabbitMQ Puppet configuration and adds the missing parameter to /etc/rabbitmq/rabbitmq-env.conf. Now rabbitmqctl does not fail in IPv6 environments

BZ#1389413

Prior to this update, HAProxy checking of MySQL resulted in a long timeout (16 seconds) before a failed node would be removed from service. Consequently, OpenStack services connected to a failed MySQL node could return API errors to users/operators/tools. With this update, the check interval settings have been reduced to drop failed MySQL nodes within 6 seconds of failure. As a result, OpenStack services should failover to working MySQL nodes much faster and produce fewer API errors to their consumers.

BZ#1262070

You can now use the director to configure Ceph RBD as a Block Storage backup target. This will allow you to deploy an overcloud where volumes are set to back up to a Ceph target. By default, volume backups will be stored in a Ceph pool called 'backups'.

Backup settings are configured in the following environment file (on the undercloud):

```
/usr/share/openstack-tripleo-heat-templates/environments/cinder-backup.yaml
```

BZ#1378391

Both Redis and RabbitMQ had a start and stop timeouts of 120s in Pacemaker. In some environments, this was not enough and caused restarts to fail. This fix increases the timeout to 200s, which is the same for the other systemd resources. Now Redis and RabbitMQ should have enough time to restart on the majority of environments.

BZ#1279554

Using the RBD backend driver (Ceph Storage) for OpenStack Compute (nova) ephemeral disks applies two additional settings to libvirt:

```
hw_disk_discard : unmap  
disk_cachemodes : network=writeback
```

This allows reclaiming of unused blocks on the Ceph pool and caching of network writes, which improves the performance for OpenStack Compute ephemeral disks using the RBD driver.

Also see <http://docs.ceph.com/docs/master/rbd/rbd-openstack/>

python-cotyledon**BZ#1374690**

Previously, a bug in an older version of `cotyledon` caused `metricsd` to not start properly and throw a traceback.

This update includes a newer 1.2.7-2 `cotyledon` package. As a result, no traceback occurs and `metricsd` starts correctly.

python-django-horizon**BZ#1198602**

This enhancement allows the `admin` user to view a list of the floating IPs allocated to instances, using the admin console. This list spans all projects in the deployment.

Previously, this information was only available from the command-line.

BZ#1328830

This update adds support for multiple theme configurations. This was added to allow a user to change a theme dynamically, using the front end. Some use-cases include the ability to toggle between a light and dark theme, or the ability to turn on a high contrast theme for accessibility reasons.

As a result, users can now choose a theme at run time.

python-django-openstack-auth

BZ#1287586

With this enhancement, domain-scoped tokens can be used to login to the Dashboard (horizon).

This was added to fully support the management of identity in keystone v3 when using a richer role set, where a domain-scoped token is required. `django_openstack_auth` must support obtaining and maintaining this type of token for the session.

As a result, horizon support for domain-scoped tokens has been available since Red Hat OpenStack Platform 9.

python-gnocchiclient

BZ#1346370

This update provides the latest client for OpenStack Telemetry Metrics (gnocchi) to support resource types.

python-ironic-lib

BZ#1381511

OpenStack Bare Metal (ironic) provides user data to new nodes through the creation of a configdrive as an extra primary partition. This requires a free primary partition available on the node's disk. However, a bug caused OpenStack Bare Metal to not distinguish between primary and extended partitions, which caused the partition count to report no free partitions available for the configdrive. This fix distinguishes between primary and extended partitions. Deployments now succeed without error.

BZ#1387148

OpenStack Bare Metal (ironic) contained parsing errors in configdrive implementation for whole disk images, which caused deployment failure. This fix corrects the return value parsing for in configdrive implementation. It is now possible to deploy whole disk images with configdrive.

-

python-tripleoclient

BZ#1364220

OpenStack Dashboard (horizon) was incorrectly included in list of services the director uses to create endpoints in OpenStack Identity (keystone). A misleading 'Skipping "horizon" postconfig' message appeared when deploying the overcloud. This fix removes horizon from the service list endpoints added to keystone and modifies the "skipping postconfig" messages to only appear in debug mode. The misleading 'Skipping "horizon" postconfig' message no longer appears.

BZ#1383930

If using DHCP HA, the ``NeutronDhcpAgentsPerNetwork`` value should be set either equal to the number of dhcp-agents, or 3 (whichever is lower), using composable roles. If this is not done, the value will default to ``ControllerCount`` which may not be optimal as there may not be enough dhcp-agents running to satisfy spawning that many DHCP servers for each network.

BZ#1384246

Node delete functions used Heat's 'parameters' instead of 'parameter_defaults'. This caused Heat to redeploy some resources, such as unintentionally redploying nodes. This fix switches the node delete functions to use only 'parameter_defaults'. Heat resources are correctly left in place and not redeployed.

python-twisted

BZ#1394150

The python-twisted package failed to install as a part of the Red Hat OpenStack Platform 10 undercloud installation due to missing "Obsoletes" for the package. This fix includes a packaging change with an "Obsoletes" list, which removes the obsolete packages during the python-twisted package installation and provides a seamless update and cleanup.

As a manual workaround, make sure not to install any python-twisted-* packages from the Red Hat Enterprise Linux 7.3 Optional repository, such as python-twisted-core. If the undercloud contains these obsolete packages, remove them with:

```
$ yum erase python-twisted-*
```

rabbitmq-server

BZ#1357522

RabbitMQ would bind to port 35672. However, port 35672 is in the ephemeral range, which leaves the possibility of other services opening up the same port. This could cause RabbitMQ to fail to start. This fix changes the RabbitMQ port to 25672, which is outside of the ephemeral port range. No other service listens on the same port and RabbitMQ starts successfully.

rhosp-release

BZ#1317669

This update includes a release file to identify the overcloud version deployed with OSP director. This gives a clear indication of the installed version and aids debugging. The overcloud-full image includes a new package (rhosp-release). Upgrades from older versions also install this RPM. All versions starting with OSP 10 will now have a release file. This only applies to Red Hat OpenStack Platform director-based installations. However, users can manually the install the rhosp-release package and achieve the same result.

sahara-image-elements

BZ#1371649

This enhancement updates the main script on `sahara-image-element` to only allow the creation of images for supported plugins. For example, you can use the following command to create a CDH 5.7 image using Red Hat Enterprise Linux 7:

```
----
>> ./diskimage-create/diskimage-create.sh -p cloudera -v 5.7
```

```
Usage: diskimage-create.sh
       [-p cloudera|mapr|ambari]
       [-v 5.5|5.7|2.3|2.4]
       [-r 5.1.0]
----
```

4.2. RHEA-2018:2670 — RED HAT OPENSTACK PLATFORM 10 ENHANCEMENT UPDATE

The bugs contained in this section are addressed by advisory RHEA-2018:2670. Further information about this advisory is available at <https://access.redhat.com/errata/RHEA-2018:2670.html>.

instack-undercloud

BZ#1582662

Previously, Undercloud upgrade failed due to dependency issues.

With this fix, the upgrade procedure automatically removes the mariadb-

devel and neutron-vpnaas packages from the Undercloud before running the FFWD. As a result, the Undercloud upgrade succeeds.

openstack-tripleo-heat-templates

BZ#1584582

NFV deployments require additional scripts to be executed as NodeUserData and NodeExtraConfigPost during the deployment. This configures the kernel args and configure tuned with a reboot before starting the puppet execution. Previously, these files were part of the documentation, but with this release, the files are in the tripleo-heat-templates repository. If the deployment has additional user changes, copy this file, add the additional changes, and use the new file for the deployment.

For OVS-DPDK deployments:

resource_registry:

OS::TripleO::Compute::NodeUserData: /usr/share/openstack-tripleo-heat-templates/firstboot/userdata_nfv_ovsdpdk.yaml

OS::TripleO::NodeExtraConfigPost: /usr/share/openstack-tripleo-heat-templates/extraconfig/post_deploy/post_deploy_nfv.yaml

For SR-IOV deployments:

resource_registry:

OS::TripleO::Compute::NodeUserData: /usr/share/openstack-tripleo-heat-templates/firstboot/userdata_nfv_sriov.yaml

OS::TripleO::NodeExtraConfigPost: /usr/share/openstack-tripleo-heat-templates/extraconfig/post_deploy/post_deploy_nfv.yaml

NOTE: These files are based only on the Compute role. If you create composable roles with NFV services, you must modify these files and registry mapping according to your requirements.

BZ#1597997

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 also be consumed by any other service.

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).
```

With this update, the new libvirtd live-migration range of 61152-61215 is not in the ephemeral range and the related failures no longer occur.

This completes the port change work started in BZ1573796.

BZ#1618797

Previously, OpenStack special handling did not upgrade all dependent packages before installing openvswitch during an upgrade, resulting in openvswitch package upgrade failure.

With this update, OpenStack special handling upgrades all dependent packages and the openvswitch upgrade is successful.

BZ#1601708

With this update, the hugetlbfs gid value correlates to the kolla fixed gid value to allow easy migration to Red Hat OpenStack Platform 13, where libvirt runs in a kolla container.

BZ#1599368

With this update, parallelization of the selinux permission change enables faster upgrade of Ceph OSD.

BZ#1599370

Previously, when upgrading from OpenStack 9 to OpenStack 10, Ceph clusters had a default status of 'health_warn'.

With this update, Ceph clusters have a status of 'OK' after upgrading to OpenStack 10.

puppet-nova**BZ#1571756**

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 that has been 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#1579699

With this enhancement, the Nova 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.

In this update, the restriction of using only the PCID flag is extended to expose multiple CPU feature flags.

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

puppet-tripleo

BZ#1605973

Previously, the xinetd service that provides the 'clustercheck' status for the Galera database was bound to all interfaces. This means that in a Controller node where a network interface was exposed to an untrusted network, an attacker could flood the xinetd service with requests and render it unavailable due to the relatively low request limit.

With this update, the clustercheck xinetd service is now bound to the network interface for internal Galera communication, and the xinetd service is no longer exposed to all network interfaces.

BZ#1520799

Previously, the cookie set in the horizon stanza for haproxy was incorrect and connections to horizon may be load balanced to the wrong server.

With this update, the cookie now points to the correct server.

BZ#1589361

Previously, it was not possible to customize the haproxy [defaults] section in the haproxy.cnf file. It was necessary to modify this section manually.

With this update, you can pass in a new hierakey and override the defaults. For example, to set the default retries to 7, you can pass in the following hierakey:

```
parameter_defaults:
  ExtraConfig:
    tripleo::haproxy::haproxy_defaults_override:
      retries: 7
```

BZ#1595315

During a version upgrade, the Block Storage Service (cinder) 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.

BZ#1598428

Previously, any non-containerized OpenStack service failed to connect to the Ceph cluster because the file ACLs mask set on the CephX keyrings blocked read permissions for non-containerized OpenStack services.

With this update, Puppet now sets the file ACLs mask for the CephX keyrings so that it can grant read permissions to specific users. This allows a non-containerized OpenStack service to connect to the Ceph cluster.

BZ#1436495

Since the release of the new HA architecture in Red Hat OpenStack Platform version 10, the majority of services are now managed by systemd. All OpenStack services are configured to restart automatically if they fail unexpectedly.

However, non-OpenStack services have a default configuration which does not enable automatic restart on failure. For example, if memcached, apache, or mongodb fail, you must restart these services manually. This may lead to service disruption if failure happens on all nodes.

With this update, the systemd unit files of these services include the option to restart automatically if the services fail.

BZ#1567368

Previously, the CinderNetappNfsMountOptions TripleO Heat parameter was inadvertently ignored, and was not used to configure the corresponding setting in the cinder Netapp backend. As a result, it was not possible to configure the Netapp NFS mount options with the TripleO Heat parameter.

With this update, the code responsible for handling the cinder Netapp configuration no longer ignores the CinderNetappNfsMountOptions parameter and the CinderNetappNfsMountOptions parameter correctly configures the cinder Netapp NFS mount options.

BZ#1598562

Previously, all overcloud nodes were deployed with the same iSCSI initiator name (IQN). This is a consequence of using a common overcloud image. Later in the deployment, the IQN is reset on Compute nodes. However, the IQN is not reset on Controllers, which also need to support iSCSI connections. As a result, all overcloud Controller nodes have the same IQN, which causes iSCSI connections to fail.

With this update, the IQN is now reset on both Controller and Compute nodes and Controllers can create reliable iSCSI connections because all of the Controllers have a unique IQN.

NOTE: The IQN on an overcloud node should be reset once, and only once. If a user has already manually reset the IQN on an overcloud node, then care must be taken to ensure that TripleO does not reset the IQN a second time.

TripleO uses a sentinel file (`/etc/iscsi/.initiator_reset`) to determine whether it should reset the node IQN. To prevent TripleO from resetting the IQN on a node, run the following command on that node:

```
sudo touch /etc/iscsi/.initiator_reset
```

python-os-brick

BZ#1572574

Previously, the OS-Brick FC code scanned all present HBAs, which could unintentionally add unwanted devices.

With this update, OS-Brick FC code scans HBAs that match the following criteria:

- HBAs in an initiator map, if present.
- HBAs that are connected in the single WWNN for all ports case.
- HBAs with wildcards.

4.3. RHEA-2018:2671 — RED HAT OPENSTACK PLATFORM 10 ENHANCEMENT UPDATE

The bugs contained in this section are addressed by advisory RHEA-2018:2671. Further information about this advisory is available at <https://access.redhat.com/errata/RHEA-2018:2671.html>.

diskimage-builder

BZ#1267169

Previously, systemd did not handle DHCP service interface names that contained '-' correctly. As a result of this, these interfaces failed to start and logged the error 'Failed to start DHCP interface'.

With this fix, systemd now escapes interface names that contain '-'.

openstack-gnocchi

BZ#1562121

Previously, Gnocchi attempted to create a storage directory on every

startup, even if the storage directory already existed. Gnocchi failed to start if the directory creation failed.

With this update, `gnocchi-upgrade` creates the storage directory only once. As a result, Gnocchi starts successfully.

openstack-ironic-python-agent

BZ#1565295

With this update, bare metal node introspection reports both the `/dev/XXX` block device name and the `/dev/disk/by-path/XXX` name. Unlike the `/dev/XXX` name, the `/dev/disk/by-path/XXX` name does not change at system reboot and may be the same across similarly configured hardware. This update improves reliability of deployments by using `/dev/disk/by-path/XXX` information in the cloud configuration.

openstack-manila

BZ#1610598

Previously, the NetApp driver ignored the size argument when a user requested a share from an existing snapshot. As a result, users did not get shares of the requested size.

With this update, the NetApp driver creates shares according to the size argument that users specify.

BZ#1610604

Previously, the NetApp driver operating in `driver_handles_share_servers=True` mode could not delete share servers that were created on non-segmented networks. This resulted in clean-up issues that prevented users from creating new shares.

With this update, the NetApp driver does not assume that share servers are provisioned only on segmented (VLAN) networks. As a result of this, share servers on non-segmented networks can be cleaned up successfully, and users can create new shares.

BZ#1610629

Previously, the security style of CIFS shares that were provisioned with the NetApp driver was incorrect. As a result of this, users were unable to write data to CIFS shares, even with explicit `'rw'` access.

With this update, the security style of the NetApp ONTAP driver is always `'ntfs'` on CIFS shares, and users that request `'rw'` access with the Shared File System service (manila) can write data to CIFS shares successfully.

BZ#1610639

Previously, the NetApp driver operating in `driver_handles_share_servers=True` mode failed to configure Active Directory services when the Active Directory server was not in the same subnet as the ONTAP Vserver. As a result, users were unable to create CIFS shares on the NetApp back-end when the Active Directory server was not on the private tenant network.

With this update, the NetApp driver creates the necessary static routes with the gateway specified on the tenant networks. Users can create CIFS shares on the NetApp back-end when the Active Directory service is on a different network, but a path exists with the tenant network gateway.

BZ#1591373

Previously, the Shared File System service (manila) emitted unhelpful warnings about ignored keywords on every wsgi request. As a result, log efficiency was reduced.

With this update, the unhelpful warnings no longer appear in logs.

BZ#1591376

Previously, the Shared File System service (manila) emitted unhelpful warnings about ignored keywords on every wsgi request. As a result, log efficiency was reduced.

With this update, the unhelpful warnings no longer appear in logs.

python-cliff**BZ#1437402**

Previously, commands with output that contained non-ASCII characters failed when the command was run with `--format=csv`.

With this update, command output that contains unicode characters displays correctly.

python-openstackclient**BZ#1576172**

Previously, the expected status for server migration completion was set only to "active". As a result, the `'openstack server migrate --wait'` command hung forever.

With this update, "verify_resize" was added to the list of expected statuses for the 'migrate' command and the `'openstack server migrate --wait'` command succeeds when the server migration is verified.

BZ#1581844

Previously, the regular expression used to match and extract the load average from Nova hypervisors load average API was incorrect. As a result, the load average would not display if the number of users was 1.

With this update, the regular expression matches all possible load average formats, and the load average displays correctly in all cases.

BZ#1563548

Previously, the command to unset static routes from a router failed with the error 'Router does not contain route xxx'.

With this fix, the command 'openstack router unset --route' is now successful.

python-os-client-config**BZ#1477126**

Previously, password values in formatted strings were expanded, causing the client commands to fail when the password contained special characters.

With this update, passwords are not subject to formatting and the client accepts passwords that contain special characters.

python-oslo-messaging**BZ#1599975**

Previously, when an OpenStack service logs at DEBUG level, Oslo Messaging logs the message "Timed out waiting for RPC response" unnecessarily.

With this fix, Oslo Messaging no longer logs this message in instances when the timeout is recoverable.

4.4. RHBA-2019:0055 — RED HAT OPENSTACK PLATFORM 10 BUG FIX UPDATE

The bugs contained in this section are addressed by advisory RHBA-2019:0055. Further information about this advisory is available at <https://access.redhat.com/errata/RHBA-2019:0055.html>.

openstack-tripleo-heat-templates**BZ#1625166**

This update fixes an OSP 9 to OSP 10 upgrade issue that sometimes

prevented the spawning of VMs during upgrades.

Prior to this update, VMs could not be spawned between ceph/compute upgrade and convergence, because ceph librados libraries were open in memory, conflicting with the upgraded client on disk. That triggered calls to non-existent (in-memory) symbols.

To work around this issue, nova-compute is restarted on compute nodes to synchronize the disk and the in-memory client libraries.

BZ#1646332

This update fixes an issue that caused OpenStack API outages and control plane loss during execution of the "pcs cluster stop" command, greatly reducing the incidence of failed requests during minor updates.

Note: In manual maintenance procedures, operators should migrate the VIPs off the affected node first.

BZ#1650702

This update fixes a configuration issue that caused failure of operations on volumes that use Nova's privileged API (for instance, migrating an in-use volume).

The failures happened because the OpenStack Platform director was not configuring authentication data required for the Block Storage service (Cinder) to access privileged portions of the Nova API.

The director now configures Cinder with Nova's authentication data. As a result, operations on volumes that require privileges succeed.

BZ#1623554

This update adds a TripleO heat template parameter as an option for setting RX/TX queue size.

Prior to this update, users could set RX/TX queue size with ``nova::compute::libvirt::rx_queue_size/nova::compute::libvirt::tx_queue_size``. However, there was no dedicated TripleO heat template parameter. With this update, users can set the RX/TX queue size either on a global level using:

```
...  
parameter_defaults:  
NovaLibvirtRxQueueSize: 1024  
NovaLibvirtTxQueueSize: 1024  
...
```

or overwrite the hieradata via `[ROLE]ExtraConfig`, which then can be used to configure a subset of compute nodes for which a dedicated role was created:

```
...
```

```
parameter_defaults:
NovaComputeExtraConfig:
  nova::compute::libvirt::rx_queue_size: 1024
  nova::compute::libvirt::tx_queue_size: 1024
  ...
```

Note: The possibilities described above are mutually exclusive.

os-net-config

BZ#1654987

This update ensures that VLAN interfaces are restarted when underlying devices are restarted after a device configuration change, allowing the successful restoration of networks.

Prior to this update, a VLAN interface was not restarted when the underlying device was restarted. Network routes using the VLAN interface as the next hop were removed and not restored.

With this update, the VLAN interfaces are restarted when the underlying devices are restarted. Network routes are restored.

puppet-tripleo

BZ#1649363

This update lets the operator specify custom timeouts for each haproxy back end via special hiera keys:

```
ExtraConfig:
tripleo::haproxy::cinder::options:
  'timeout client': '90m'
  'timeout server': '90m'
```

With this support an operator can specify custom options for each haproxy back end.

python-os-brick

BZ#1599641

This update fixes an issue that sometimes prevented detachment of multipath devices.

Prior to this update, detachment of multipath devices included a flush of each individual path. If an individual path flushing failed, the detachment failed, even though there were other paths available to flush all the data.

Because flushing the multipath already ensures buffered data is written

on the remote device, individual paths are no longer flushed. As a result, detaching only fails when the detachment would actually result in data loss.

BZ#1583466

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#1634163

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.