Red Hat OpenStack Platform 8

Release Notes

Release details for Red Hat OpenStack Platform 8

Red Hat OpenStack Platform 8 Release Notes

Release details for Red Hat OpenStack Platform 8

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

Legal Notice

Abstract

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

Table of Contents

CHAPTER 1. INTRODUCTION

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 Liberty, 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 "Liberty" 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 "Liberty" release itself are available at the following location: https://wiki.openstack.org/wiki/ReleaseNotes/Liberty

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

This version of Red Hat OpenStack Platform is supported on Red Hat Enterprise Linux 7.2 or later.

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 recommended best practices for installing 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. 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.

1.8. CONTENT DELIVERY NETWORK (CDN) CHANNELS

This section describes the channel and repository settings required to deploy Red Hat OpenStack Platform 8.

You can install Red Hat OpenStack Platform 8 through the Content Delivery Network (CDN). To do so, configure **subscription-manager** to use the correct channels.

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 8 for RHEL 7 (RPMs)	`rhel-7-server-openstack-8-rpms`
Red Hat OpenStack Platform 8 director for RHEL 7 (RPMs)	`rhel-7-server-openstack-8-director-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	`rhel-7-server-optional-rpms`
Red Hat OpenStack Platform 8 Operational Tools for RHEL 7 (RPMs)	`rhel-7-server-openstack-8-optools-rpms`

Channels to Disable

The following table outlines the channels you must disable to ensure Red Hat OpenStack Platform 8 functions correctly.

Table 1.3. Channels to Disable

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

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

1.9. 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

Red Hat OpenStack Platform 8 brings several notable new enhancements for the director:

- Broader support for Cisco Networking through Neutron, including:

  - N1KV ML2 plug-in

  - N1KV VEM and VSM modules

  - Nexus 9K ML2 plug-in

  - UCSM ML2 plug-in

- New parameters for network configuration in environment files, including **type_drivers**, **service_plugins** and **core_plugin**.

- Big Switch Networks support, including Big Switch ML2 plugin, LLDP, and bonding support

- VXLAN is now the default overlay network. This is because VXLAN performs better and NICs with VXLAN offload are more common.

- MariaDB's maximum number of connections now scales with the number of CPU cores in the Controller nodes.

- The director can now set RabbitMQ's file descriptor limits.

- SSL support for Red Hat OpenStack Platform components deployed on nodes in overclouds.

- IPv6 support for overcloud nodes.

2.2. BLOCK STORAGE

The following section briefly describes the new features included in the Block Storage service for Red Hat OpenStack Platform 8.

Generic Volume Migration

Generic Volume Migration allows volume drivers that do not support iSCSI and use other means for data transport to participate in volume migration operations. It uses *create_export* to create and attach a volume via iSCSI to perform I/O operations. By making this more generic, we can allow other drivers to take part in volume migration as well.

This change is necessary to support volume migration for the Ceph driver.

Import/Export snapshots

Provides a means to import and export snapshots. Import/export snapshots function is a complement for the import/export volume functionality.

- It provides the ability to import volumes' snapshot from one Block Storage volume to another, and import *non* OpenStack snapshots already on a back-end device into OpenStack Block Storage service.

- Export snapshots works the same way as export volumes.

Non Disruptive Backup

Previously, a backup operation could only be performed when a volume was detached. You can now back up a volume by using the following steps:

- Take a temporary snapshot

- Attach snapshot

- Do backup from the snapshot

- Cleanup temporary snapshot

For an attached volume, taking a temporary snapshot is usually less expensive than creating a whole temporary volume. You can now attach the snapshot and read it directly.

If a driver has not implemented attach snapshot and does not have a way to read from a snapshot, you can create a temporary volume from the attached source volume, and backup the temporary volume.

New Volume Replication API

Volume replication is a key storage feature and a requirement for features such as high-availability and disaster recovery of applications running on OpenStack clouds. This release adds initial support for volume replication in the Block Storage service, and includes support for:

- Replicate volumes (primary to secondary approach)

- Promote a secondary to primary (and stop replication)

- Re-enable replication

- Test that replication is running properly

Generic Image Cache

Currently some volume drivers implement the *clone_image* method and use an internal cache of volumes on the backend that hold recently used images. For storage backends that can do very efficient volume clones it is a potentially large performance improvement over having to attach and copy the image contents to a each volume. To make this functionality easier for other volume drivers to use, and prevent any duplication in the code base, image cache is added.

Use this functionality when creating a volume from an image more than once. As an end user, you will see (potentially) faster volume creation from an image after the first time.

2.3. COMPUTE

Red Hat OpenStack Platform 8 brings several notable new features in the Compute Service:

- The **nova set-password** *server* command, which changes the admin password for a server, is now available.

- The libvirt driver has been enhanced to enable *virtio-net multiqueue* for instances. With this feature on, workload is scaled across vCPUs, thereby increasing network performance.

- Disk QoS (Quality of Service) when using Ceph RBD (RADOS block device) storage. Among other things, sequential read or write limitation, and total allowed IOPS or bandwidth for a guest can be configured.

- The *Mark host down* API for external high-availability solutions. This API allows external tools to notify the Compute service of compute node failure, which improves instance resiliency.

2.4. IDENTITY

Red Hat OpenStack Platform 8 introduces a number of new features for Identity Service:

- You can now configure Identity Provider-specific WebSSO. Previously, you had to configure WebSSO globally for *keystone*. With this update, you can configure WebSSO for each Identity Provider, directing dashboard queries to the individual endpoints, rather than performing additional discovery steps.

- New attributes are available for SAML assertion: `openstack_user_domain` for mapping user domains, and `openstack_project_domain` for mapping project domains.

- Experimental support has been added for *keystone* tokenless authorization using X.509 SSL client certificates.

2.5. IMAGE SERVICE

The following section briefly describes the new features included in the Image service for Red Hat OpenStack Platform 8.

Image Signing and Encryption

This feature provides support for image signatures and signature verification, allowing the user to verify that an image has not been modified prior to booting the image.

Artifacts Repository (experimental API)

This feature extends Image service functionality to store not only the virtual machine images but any other artifacts, such as, binary objects accompanied with composite metadata.

Image service becomes a catalog of such artifacts, providing capabilities to store, search and retrieve artifacts, their metadata and associated binary objects.

2.6. OBJECT STORAGE

This release also includes a new ring tool, namely the **Ring Builder Analyzer**. It is used for analyzing how well the ring builder performs its job in a particular scenario.

The ring builder analyzer takes a scenario file containing some initial parameters for a ring builder plus a certain number of rounds. In each round, some modifications are made to the builder, e.g. add a device, remove a device, change a device's weight. Then, the builder is repeatedly rebalanced until it settles down. Data about that round is printed, and the next round begins.

2.7. OPENSTACK NETWORKING

2.7.1. QoS

Red Hat OpenStack Platform 8 introduces support for network quality-of-service (QoS) policies. These policies allow OpenStack administrators to offer varying service levels by applying rate limits to ingress and egress traffic for instances. Any traffic that exceeds the specified rate is consequently dropped.

2.7.2. Open vSwitch Update

Open vSwitch (OVS) has been updated to the upstream 2.4.0 release. This update includes a number of notable enhancements:

- Support for the Rapid Spanning Tree Protocol (IEEE 802.1D-2004), allowing faster convergence after topology changes.

- Optimized multicast efficiency with support for IP multicast snooping (IGMPv1, IGMPv2 and IGMPv3).

- Support for vhost-user, a QEMU feature that offers improved I/O efficiency between a guest and a user-space vSwitch.

- OVS version 2.4.0 also includes various performance and stability improvements.

For further details on Open vSwitch 2.4.0, see http://openvswitch.org/releases/NEWS-2.4.0

2.7.3. RBAC for Networks

Role-based Access Control (RBAC) policies in OpenStack Networking allows granular control over shared *neutron* networks. Previously, networks were shared either with all tenants, or not at all. OpenStack Networking now uses a RBAC table to control sharing of *neutron* networks between tenants, allowing an administrator to control which tenants are granted permission to attach instances to a network.
As a result, cloud administrators can remove the ability for some tenants to create networks, and can instead allow them to attach to pre-existing networks that correspond to their project.

2.8. TECHNOLOGY PREVIEWS

NOTE

For more information on the support scope for features marked as technology previews, see https://access.redhat.com/support/offerings/techpreview/.

2.8.1. New Technology Previews

The following new features are provided 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.

DPDK-Accelerated Open vSwitch

The Data Plane Development Kit (DPDK) consists of a set of libraries and user-space drivers for fast packet processing, enabling applications to perform their own packet processing directly to/from the NIC, delivering up to wire speed performance for certain use cases. In addition, OVS+DPDK significantly improves the performance of Open vSwitch while maintaining its core functionality. It enables the packet switching from the host's physical NIC to the application in the guest instance (and between guest instances) to be handled almost entirely in user-space.

In this release, the OpenStack Networking (neutron) OVS plugin was updated to support OVS+DPDK back end configuration. OpenStack projects can now use the neutron API to provision networks, subnets and other networking constructs, while using OVS+DPDK to gain improved network performance for instances.

OpenDaylight Integration

Red Hat OpenStack Platform 8 now 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 8 is limited to the modules required to support OpenStack deployments using OVSDB NetVirt, and is based on the upstream Beryllium version. The following packages provide the Technology Preview: opendaylight, networking-odl

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.

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.

2.8.2. Previously Released Technology Previews

The following features remain as technology previews:

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 Enterprise Linux OpenStack Platform also provides fully supported methods for dividing compute resources in Red Hat Enterprise Linux OpenStack Platform; namely, Regions, Availability Zones, and Host Aggregates. For more information, see Manage Host Aggregates.

Database-as-a-Service (DBaaS)

OpenStack Database-as-a-Service allows users to easily provision single-tenant databases within OpenStack Compute instances. The Database-as-a-Service framework allows users to bypass much of the traditional administrative overhead involved in deploying, using, managing, monitoring, and scaling databases.

Distributed Virtual Routing

Distributed Virtual Routing (DVR) allows you to place L3 Routers directly on Compute nodes. As a result, instance traffic is directed between the Compute nodes (East-West) without first requiring routing through a Network node. Instances without floating IP addresses still route SNAT traffic through the Network node.

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.

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.

File Share Service

The OpenStack File Share Service provides a seamless and easy way to provision and manage shared file systems in OpenStack. These shared file systems can then be used (mounted) securely to instances. The File Share Service also allows for robust administration of provisioned shares, providing the means to set quotas, configure access, create snapshots, and perform other useful admin tasks.

The following section briefly describes the new features included in the File Share Service for Red Hat OpenStack Platform 8.

Manila Horizon dashboard plug-in

With this release, users can now interact with the capabilities that the File Share Service provides via the dashboard, including an interactive menu for creating and working with shares.

Share migration

Share migration is a new feature that enables migration of shares from back end to back end.

The following approaches are available:

- Delegate to driver - This is a very optimized but restricted approach. The driver is able to perform migration in a more efficient way if it understands the destination backend. A model update should be returned by the driver after the migration.

- Manages coordinate, delegates some tasks to drivers - This approach creates a new share on destination host, mount both exports from the manila node, copy all files and then delete the old share. This approach should work for any driver that implements some methods necessary to help the migration process, such as:

  - changing the source share to read-only so users are less impacted by migration.

  - mounting/unmounting exports with specific protocols.

For the second to work, every driver during server_setup method must create a port that allows connectivity between the share server and manila node.

Availability zones

The File Share Service client's share creation code now accepts and uses availability zone arguments. This also allows the preservation of availability zone information when creating a share from a snapshot.

Oversubscription in thin provisioning

This release adds support for oversubscription in thin provisioning, which addresses the use case where certain drivers still report *infinite* or *unknown* for their capacity, potentially leading to oversubscription. This update adds the following parameters:

- **max_over_subscription_ratio**: A floating-point number that represents the ratio of over subscription to be applied. This ratio is calculated as the ratio of provisioned storage to total available capacity. As such, an over subscription ratio of `1.0` means that the total amount of provisioned storage cannot exceed the total amount of available storage, whereas an over subscription ratio of 2.0 means that the total amount of provisioned storage can reach double the total amount of available storage.

- **provisioned_capacity**: The apparent amount of storage that has been provisioned. The value of this paramter is used to calculate the **max_over_subscroption_ratio**.

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.

Operational Tools

Operational Tools are logging and monitoring tools which facilitate troubleshooting. With a centralized, easy-to-use analytics and search dashboard, troubleshooting is simplified, and features such as service availability checking, threshold alarm management, and collecting and presenting data using graphs are available.

VPN-as-a-Service (VPNaaS)

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

Time-Series-Database-as-a-Service (TSDaaS)

Time-Series-Database-as-a-Service (**gnocchi**) is a multi-tenant, metrics and resource database. It is designed to store metrics at a very large scale while providing access to metrics and resources information to operators and users.

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.

3.1. ENHANCEMENTS

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

BZ#978365

> The ability of the libvirt driver to set the admin password has been added. To use this feature, run the following command: "nova root-password [server]".

BZ#978365

> The ability of the libvirt driver to set the admin password has been added. To use this feature, run the following command: "nova root-password [server]".

BZ#1042947

> This update adds support for volume migrations of the Block Storage (cinder) service. These are done in the 'Volumes' panel of the OpenStack dashboard (Project-> Compute -> Volumes and in Admin-> System Panel-> Volumes). You can perform this action on the 'Volumes' row in the table. The final patch in this series resolved the command action itself; it had previously errored out due to incorrect parameters, and parameter count issues.

BZ#1100542

> OpenStack dashboard tables summarize information about a large number of entities. This update adds a table enhancement that enables this information to be displayed within the table as a slide-down "drawer" that is activated when you click on a toggle switch within a row. The drawer appears as an additional row (with configurable height) and contains additional information about the entity in the row above it (e.g. additional entity details, metrics, graphs, etc.). Multiple drawers may be opened at one time.

BZ#1104445

> Instances can now be cold migrated or live migrated from hosts marked for maintenance. A new action button in the System > Hypervisors > Compute Host tab in the dashboard allows administrative users to set options for instance migration.
>
> Cold migration moves an instance from one host to another, reboots across the move, and its destination is chosen by the scheduler. This type of migration should be used when the administrative user did not

select the 'live_migrate' option in the dashboard or the migrated
instance is not running.

Live migration moves an instance (with "Power state" = "active") from
one host to another, the instance doesn't appear to reboot, and its
destination is optional (it can be defined by the administrative user or
chosen by the scheduler). This type of migration should be used when the
administrative user selected the 'live_migrate' option in the dashboard
and the migrated instance is still running.

BZ#1149599

With this feature, you can now use Block Storage (cinder) to create a
volume by specifying either the image ID or image name.

BZ#1166963

This update replaces the network topology with curvature based graph as
the previous UI did not work well with larger number of nodes or
networks.

The new network topology map can handle more nodes, looks stylish and
the node layout can be re-organized.

BZ#1167563

The 'Launch Instance' workflow has been redesigned and re-implemented to
be more responsive with this update.

1. To enable this update, add the following values in your
/etc/openstack-dashboard/local_settings file:

LAUNCH_INSTANCE_LEGACY_ENABLED = False
LAUNCH_INSTANCE_NG_ENABLED = True

2. Restart 'httpd':
systemctl restart httpd

BZ#1167565

This update adds a common API hosted by the Image Service (glance) for
vendors, admins, services, and users to meaningfully define an available
key/value pair, and tag metadata. The intent is to enable better
metadata collaboration across artifacts, services, and projects for
OpenStack users.
This definition describes the available metadata that can be used on
different types of resources (images, artifacts, volumes, flavors,
aggregates, among others). A definition includes the properties type,
key, description, and constraints. This catalog will not store the
values for specific instance properties.
For example, a definition of a virtual CPU topology property for a
number of cores will include the key to use, a description, and value
constraints, such as requiring it to be an integer. As a result, users
(potentially through the dashboard) would be able to search this catalog

to list the available properties they can add to a flavor or image. They will see the virtual CPU topology property in the list and know that it must be an integer. In the dashboard example, when the user adds the property, its key and value will be stored in the service that owns that resource (in nova for flavors, and in glance for images).

BZ#1168359

Nova's serial console API is now exposed for instances. Specifically, a serial console is available for hypervisors not supporting VNC or Spice. This update adds support for it in the dashboard.

BZ#1189502

With this update, configuration settings now exist to set timeouts, after which clusters which have failed to reach the 'Active' state will be automatically deleted.

BZ#1189517

When creating a job template intended for re-use, you can now register a variable for datasource URLs with OpenStack Data Processing (sahara). Doing so allows you to easily change input and output paths per run, rather than an actual URL (which would require revising the template, or manually revising the URL per run between jobs).

This makes it easier to reuse job templates when data source jobs are mutable between runs, as is true for most real-world cases.

BZ#1192641

With this release, in order to provide security isolation, the '/usr/local' path has been removed from the default Block Storage rootwrap configuration. As a result, the deployments relying on Block Storage service executing commands from the '/usr/local/' as the 'root' user will need to add configuration for the commands to work.

BZ#1212158

This updates provides OpenStack notifications. Previously there were external consumers of OpenStack notifications that could not interface with director-deployed cloud because notifications were not enabled. Now director enables notifications for external consumers.

BZ#1214230

With this update, a new feature adds pagination support for the Block Storage 'snapshots-list' and 'backups-list' commands. You can now limit, marker and sort parameters to control the number of returned results, starting element and their order.

Retrieving a limited number of results instead of the entire data set can be extremely useful on the large deployments with thousands of

snapshots and backups.

BZ#1225163

The Director now properly enabled notifications for external consumers.

BZ#1229634

Previously, there was no secure way to remotely access S3 backend in a private network.

With this update, a new feature allows Image service S3 driver to connect a S3 backend from a different network in a secure way through the HTTP proxy.

BZ#1238807

This enhancement enables the distribution of per-node hieradata, matching the nodes from their UUID (as reported by 'dmidecode'). This allows you to scale CephStorage across nodes equipped with a different number/type of disks.
As a result, CephStorage nodes can now be configured with non-homogeneous disk topologies. This is done by provisioning a different configuration hash for the ceph::profile::params::osds parameter.

BZ#1249601

OpenStack Bare Metal (ironic) now supports deploying nodes in UEFI mode. This is due to requests from customers with servers that only support UEFI boot.

BZ#1257306

This release includes a tech preview of Image Signing and Verification for glance images. This feature helps protect image integrity by ensuring no modifications occur after the image is uploaded by a user. This capability includes both signing of the image, and signature validation of bootable images when used.

BZ#1258643

To provide better flexibility for administrators on deployments with an assortment of storage backends, Block Storage now defines standard names for the capabilities, for example, QoS, compression, replication, bandwidth control, and thin provisioning. This means volume type specifications that will work with multiple drivers without modifications can be defined.

BZ#1258645

This enhancement adds a new scaled backend replication implementation (between backends) that leaves the bulk of the work up to the driver,

while providing basic admin API methods. This is available where
replication is set at the volume types level, and when the cinder driver
reports its capabilities. New configuration options are available:
replication_enabled - set to True
replication_type - async, sync
replication_count - Number of replicas

BZ#1259003

The domain name for overcloud nodes defaulted to 'localdomain'. For
example: 'overcloud-compute-0.localdomain'. This enhancement provides a
parameter (CloudDomain) to customize the domain name. Create an
environment file with the CloudDomain parameter included in the
'parameter_defaults" section. If no domain name is defined, the Heat
templates default to 'localdomain'.

BZ#1262106

This enhancement enables backup of Block Storage (cinder) volumes to a
Ceph object store using the same user interface as that for backing up
cinder volumes to Object Storage (swift).
This was done to avoid the need for a second object store if Ceph was
already being used.

BZ#1266104

This update adds neutron QoS (Quality of Service) extensions to provide
better control over tenant networking qualities and limits. Overclouds
are now deployed with Neutron QoS extension enabled.

BZ#1266156

The OpenDaylight OpenStack neutron driver has been split from the
neutron project and moved to a new package, python-networking-odl.
Operators still have the driver available for use as part of their Red
Hat OpenStack Platform installations.

BZ#1266219

The Director can now deploy the Block Storage service with a Dell
EqualLogic or Dell Storage Center appliance as a back end. For more
information, see:

https://access.redhat.com/documentation/en/red-hat-openstack-
platform/version-8/dell-equallogic-back-end-guide/
https://access.redhat.com/documentation/en/red-hat-openstack-
platform/8/dell-storage-center-back-end-guide/dell-storage-center-back-
end-guide

BZ#1267951

This update introduces nested quotas. Deployers now have the ability to
manage a hierarchy of quotas in Cinder, with subprojects that inherit

> from parent projects.

BZ#1273303

> The Director now supports the OpenStack Networking
> 'enable_isolated_metadata' option. This option allows access to
> instance metadata on VMs on external routers or on isolated networks.

BZ#1279812

> With this release, panels are configurable. You can add or remove panels
> by using configuration snippets.
>
> For example, to remove the "Resource panel":
>
> * Place a file in '/usr/share/openstack-
> dashboard/openstack_dashboard/local/enabled'.
> * Name that file '_99_disable_metering_dashboard.py'.
> * Copy the following content into the file:
>
> # The slug of the panel to be added to HORIZON_CONFIG. Required.
> PANEL = 'metering'
> # The slug of the dashboard the PANEL associated with. Required.
> PANEL_DASHBOARD = 'admin'
> # The slug of the panel group the PANEL is associated with.
> PANEL_GROUP = 'admin'
> REMOVE_PANEL = True
>
> * Restart the Dashboard httpd service:
> # systemctl restart httpd
>
> For more information, see the Pluggable Dashboard Settings section in
> the Configuration Reference Guide in the Red Hat OpenStack Platform
> Documentation Suite available at:
> https://access.redhat.com/documentation/en/red-hat-enterprise-linux-
> openstack-platform/

BZ#1282429

> This update adds new parameters to configure API worker process counts,
> which allows you to tune overcloud's memory utilization and request
> processing capacity. The parameters are: CeilometerWorkers,
> CinderWorkers, GlanceWorkers, HeatWorkers, KeystoneWorkers,
> NeutronWorkers, NovaWorkers, and SwiftWorkers.

BZ#1295690

> Previously, a router that was neither an HA nor a DVR router could not
> be converted into an HA router. Instead, it was necessary to create a
> new router and reconnect all the resources (interfaces, networks etc.)
> from the old router to the new one. This update adds the ability to
> convert a legacy router into an HA or non-HA router in a few simple
> commands:

```
# neutron router-update ROUTER --admin-state-up=False
# neutron router-update ROUTER --ha=True/False
# neutron router-upgrade ROUTER --admin-state-up=True

Replace ROUTER with the ID or name of the router to convert.
```

BZ#1296568

```
With this update, overcloud nodes can now be registered to Satellite 5
with the Red Hat OpenStack Platform director, and Satellite 5 can now
provide package updates to overcloud nodes.

To register overcloud nodes with a Satellite 5 instance, pass the
following options to the "openstack overcloud deploy" command:

    --rhel-reg
    --reg-method satellite
    --reg-org 1
    --reg-force
    --reg-sat-url https://<satellite-5-hostname>
    --reg-activation-key <satellite-5-activation-key>
```

BZ#1298247

```
The Director now supports new parameters that control whether to disable
or enable the following OpenStack Networking services:

* dhcp_agent
* l3_agent
* ovs_agent
* metadata_agent

This enhancement allows the deployment of Neutron plug-ins that replace
any of these services. To disable all of these services, use the
following parameters in your environment file:

  NeutronEnableDHCPAgent: false
  NeutronEnableL3Agent: false
  NeutronEnableMetadataAgent: false
  NeutronEnableOVSAgent: false
```

BZ#1305023

```
This update allows the Dashboard (horizon) to accept an IPv6 address as
a VIP address to a Load Balancing Pool. As a result, you can now use
Dashboard to configure IPv6 addresses on a Load Balancing Pool.
```

BZ#1312373

```
This update adds options to configure Ceilometer to store events, which
can be retrieved later through Ceilometer APIs. This is an alternative
to listening to the message bus to capture events. A brief outline of
the configuration is in https://bugzilla.redhat.com/show_bug.cgi?
id=1318397.
```

—

BZ#1316235

> With the Red Hat OpenStack Platform 8 release, the inbuilt
> implementation of Amazon EC2 API in the OpenStack Compute (nova) service
> is deprecated and will be removed in the future releases.
>
> Moving forward, with the Red Hat OpenStack Platform 9 release, a new
> standalone EC2 API service will be available.

BZ#1340717

> This update removes unnecessary downtime caused by updating OvS switch
> reconfiguration when restarting the OvS agent. Previously, dropping
> flows on physical bridges caused networking to drop. The same issue was
> experienced when the patch port between br-int and br-tun was deleted
> and rebuilt during startup. This enhancement resolves these issues,
> making it possible to restart the OvS agent without unnecessarily
> disrupting network traffic. This results in no downtime when restarting
> the OvS neutron agent if the bridge is already set up and
> reconfiguration was not requested.

3.2. TECHNOLOGY PREVIEW

The items listed in this section are provided as Technology Previews. For further information on the scope of Technology Preview status, and the associated support implications, refer to https://access.redhat.com/support/offerings/techpreview/.

BZ#1322944

> This update provides the following technology preview:
>
> The director provides an option 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.

3.3. RELEASE NOTES

This section outlines important details about the release, including recommended practices and notable changes to Red Hat OpenStack Platform. You must take this information into account to ensure the best possible outcomes for your deployment.

BZ#1244555

> When the Block Storage service creates volumes from images, it downloads
> images from the Image service into
> an image conversion directory. This directory is defined by the
> 'image_conversion_dir' option in /etc/cinder/cinder.conf (under the
> [DEFAULT] section). By default, 'image_conversion_dir' is set to
> /var/lib/cinder/conversion.

If the image conversion directory runs out of space (typically, if multiple volumes are created from large images simultaneously), any attempts to create volumes from images will fail. Further, any attempts to launch instances which would require the creation of volumes from images will fail as well. These failures will continue until the image conversion directory has enough free space.

As such, you should ensure that the image conversion directory has enough space for the typical number of volumes that users simultaneously create from images. If you need to define a non-default image conversion directory, run:

```
    # openstack-config --set /etc/cinder/cinder.conf DEFAULT
image_conversion_dir <NEWDIR>
```

Replace <NEWDIR> with the new directory. Afterwards, restart the Block Storage service to apply the new setting:

```
    # openstack-service restart cinder
```

BZ#1266050

The Open vSwitch (openvswitch) package is now re-based to upstream version 2.4.0.

BZ#1300735

With this release, the 'Metering' panel in Dashboard (horizon) has been disabled due to performance issues.

3.4. KNOWN ISSUES

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

BZ#1177611

A known issue has been identified for interactions between High Availability (VRRP) routers and L2 Population. Currently, when connecting a HA router to a subnet, HA routers use a distributed port by design. Each router has the same port details on each node that it's scheduled on, and only the master router has IPs configured on that port; all the slaves have the port without any IPs configured. Consequently, L2Population uses the stale information to advise that the router is present on the node (which it states in the port binding information for that port).
As a result, each node that has a port on that logical network has a tunnel created only to the node where the port is presumably bound. In addition, a forwarding entry is set so that any traffic to that port is sent through the created tunnel.
However, this action may not succeed as there is not guarantee that the master router is on the node specified in the port binding. Furthermore,

in the event that the master router is in fact on the node, a failover event would cause it to migrate to another node and result in a loss of connectivity with the router.

BZ#1234601

The Ramdisk and Kernel images booted without specifying a particular interface. This meant the system booted from any network adapter, which caused problems when more than one interface was on the Provisioning network. In those cases it was necessary to specify which interface the system should use to boot. The interface specified should correspond to the interface which carried the MAC address from the instackenv.json file.

As a workaround, copy and paste the following block of text as the root user into the director's terminal.This creates a systemd startup script sets these parameters on every boot.

The script contains a sed command which includes "net0/mac". This sets the director to use the first Ethernet interface. Change this to "net1/mac" to use the second interface, and so on.

```
#######################################
cat << EOF > /usr/bin/bootif-fix
#!/usr/bin/env bash

while true;
        do find /httpboot/ -type f ! -iname "kernel" ! -iname "ramdisk"
! -iname "*.kernel" ! -iname "*.ramdisk" -exec sed -i 's|{mac|
{net0/mac|g' {} +;
done
EOF

chmod a+x /usr/bin/bootif-fix

cat << EOF > /usr/lib/systemd/system/bootif-fix.service
[Unit]
Description=Automated fix for incorrect iPXE BOOFIF

[Service]
Type=simple
ExecStart=/usr/bin/bootif-fix

[Install]
WantedBy=multi-user.target
EOF

systemctl daemon-reload
systemctl enable bootif-fix
systemctl start bootif-fix

#######################################
```

The bootif-fix script runs on every boot. This enables booting from a

specified NIC when more than one NIC is on the Provisioning network. To
disable the service and return to the previous behavior, run "systemctl
disable bootif-fix" and reboot.

BZ#1237009

The swift proxy port is denied in the Undercloud firewall. This means
the swift proxy only accepts connections from localhost. As a
workaround, open the swift proxy port in the firewall:

sudo iptables -I INPUT -p tcp --dport 8080 -j ACCEPT

This enabled connections to the swift proxy from remote machines.

BZ#1268426

There is a known issue that can occur when IP address conflicts are
identified on the Provisioning network. As a consequence, discovery
and/or deployment tasks will fail for hosts which are assigned an IP
address which is already in use.
You can work around this issue by performing a port scan of the
Provisioning network. Run from the Undercloud node, this will help
validate whether the IP addresses used for the discovery and host IP
ranges are available for allocation. You can perform this scan using the
nmap utility. For example (replace the network with the subnet of the
Provisioning network (in CIDR format)):

$ sudo yum install -y nmap
$ nmap -sn 192.0.2.0/24

As a result, if any of the IP addresses in use will conflict with the IP
ranges in undercloud.conf, you will need to either change the IP ranges,
or free up the IP addresses before running the introspection process, or
deploying the Overcloud nodes.

BZ#1272591

The Undercloud used the Public API to configure service endpoints during
the post-deployment stage. This meant the Undercloud needed to reach the
Public API in order to complete the deployment. If the External uplink
on the Undercloud is not the same subnet as the Public API, the
Undercloud requires a route to the Public API and any firewall ACLs must
allow this traffic. With this route, the Undercloud connects to the
Public API and completes post-deployment tasks.

BZ#1290881

The default driver with Block Storage service is the internal LVM
software iSCSI driver. This is the volume back-end which manages local
volumes.

However, the Cinder iSCSI LVM driver has significant performance issues.
In production environments, with high I/O activity, there are many
potential issues which could affect performance or data integrity,

Red Hat strongly recommends using a certified Block Storage plug-in
provider for storage in a production environment. The software iSCSI LVM
driver should be used and is only supported for single node evaluations
and proof of concept environments.

BZ#1293979

Updating packages on the Undercloud left the Undercloud in an
indeterminate state. This meant some Undercloud services were disabled
after the package update and could not start again. As a workaround, run
'openstack undercloud install' to reconfigure all Undercloud services.
After the command complete, the Undercloud services operate normally.

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#1463061

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.

3.5. DEPRECATED FUNCTIONALITY

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

BZ#1295573

With Red Hat OpenStack Platform 8 (Liberty), Red Hat develops a tighter

integration with Ceph as an end-to-end storage solution. All future
support efforts will be directed accordingly.

Presently, Red Hat Ceph Storage is fully supported as a back end for
both Block Storage and Object Storage services. In the coming major
releases, Ceph will be fully supported as a back end for ALL storage-
consuming OpenStack components. This will provide a unified storage
solution for users who wish to use commodity storage hardware.

In line with this, Red Hat Gluster Storage is deprecated as of this
release, and all Gluster-related drivers will be removed in a future
release. If you wish to continue using commodity storage hardware for
future updates, you should migrate all Gluster-backed services
accordingly.

Further, the usage of Red Hat Gluster Storage with the File Share
Service (Manila) will not be supported in this or any future release.
The respective Gluster drivers for this will also be removed.

Red Hat OpenStack Platform will continue to support the use of GlusterFS
volumes as a back end for the Image service (SwiftOnFile).

BZ#1296135

With this release, support for PowerDNS (pdns) has been removed due to a
known security issue with PolarSSL/mbedtls.

Designate can now be used with BIND9 as a backend.

BZ#1312889

This update removes the Tuskar API service from Red Hat OpenStack
Platform director 8. Tuskar was installed and configured on the
Undercloud, including an endpoint existing in the Keystone service
catalog. The RPM is no longer installed, the service is not configured,
and the endpoint is not created in the service catalog.

CHAPTER 4. TECHNICAL NOTES

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

4.1. RHEA-2016:0603 - RED HAT OPENSTACK PLATFORM 8 ENHANCEMENT ADVISORY

The bugs contained in this section are addressed by advisory RHEA-2016:0603. Further information about this advisory is available at https://access.redhat.com/errata/RHEA-2016:0603.html.

diskimage-builder

BZ#1307001

> The diskimage-builder package has been upgraded to upstream version
> 1.10.0, which provides a number of bug fixes and enhancements over the
> previous version. Notably, the python-devel package is no longer removed
> by default, as it previously caused other packages to be removed as
> well.

memcached

BZ#1299075

> Previously, memcached was unable to bind IPv6 addresses, resulting in
> memcached failing to start in IPv6 environments.
> This update addresses this issue, with memcached-1.4.15-9.1.el7ost now
> IPv6-enabled.

mongodb

BZ#1308855

> This rebase package adds improved performance for range queries.
> Specifically, queries that used the `$or` operator were previously
> affected with the 2.4 release. Those regressions are now fixed in 2.6

openstack-cinder

BZ#1272572

> Previously, a bug in the Block Storage component caused it to be
> incompatible with the Identity API v2 when working with quotas,
> resulting in failures when managing information on quotas in Block
> Storage. With this update, Block Storage has now been updated to be
> compatible with the Identity API v2, and the dashboard can now correctly
> retrieve information on volume quotas.

BZ#1295576

Previously, a bug in the cinder API server quota code used
`encryption_auth_url` when it should have used `auth_uri`.
Consequently, cinder failed to talk to keystone when querying quota
information, causing the client to receive HTTP 500 errors from cinder.
This issue has been resolved in 7.0.1.
Fix: Fixed in Cinder API service in 7.0.1, resulting in expected
behavior of the cinder quota commands.

BZ#1262106

This enhancement enables backup of Block Storage (cinder) volumes to a
Ceph object store using the same user interface as that for backing up
cinder volumes to Object Storage (swift).
This was done to avoid the need for a second object store if Ceph was
already being used.

BZ#1179445

Previously, when Ceph was used as the backing store for Block Storage
(cinder), operations such as deleting or flattening a large volume may
have blocked other driver threads.
Consequently, deleting and flattening threads may have prevented cinder
from doing other work until they completed.
This fix changes the delete and flattening threads to run in a sub-
process, rather than as green threads in the same process.
As a result, delete and flattening operations are run in the background
so that other cinder operations (such as volume creates and attaches)
can run concurrently.

BZ#1192641

With this release, in order to provide security isolation, the
'/usr/local' path has been removed from the default Block Storage
rootwrap configuration. As a result, the deployments relying on Block
Storage service executing commands from the '/usr/local/' as the 'root'
user will need to add configuration for the commands to work.

BZ#1258645

This enhancement adds a new scaled backend replication implementation
(between backends) that leaves the bulk of the work up to the driver,
while providing basic admin API methods. This is available where
replication is set at the volume types level, and when the cinder driver
reports its capabilities. New configuration options are available:
replication_enabled - set to True
replication_type - async, sync
replication_count - Number of replicas

BZ#1258643

> To provide better flexibility for administrators on deployments with an assortment of storage backends, Block Storage now defines standard names for the capabilities, for example, QoS, compression, replication, bandwidth control, and thin provisioning. This means volume type specifications that will work with multiple drivers without modifications can be defined.

BZ#1267951

> This update introduces nested quotas. Deployers now have the ability to manage a hierarchy of quotas in Cinder, with subprojects that inherit from parent projects.

openstack-glance

BZ#1167565

> This update adds a common API hosted by the Image Service (glance) for vendors, admins, services, and users to meaningfully define an available key/value pair, and tag metadata. The intent is to enable better metadata collaboration across artifacts, services, and projects for OpenStack users.
> This definition describes the available metadata that can be used on different types of resources (images, artifacts, volumes, flavors, aggregates, among others). A definition includes the properties type, key, description, and constraints. This catalog will not store the values for specific instance properties.
> For example, a definition of a virtual CPU topology property for a number of cores will include the key to use, a description, and value constraints, such as requiring it to be an integer. As a result, users (potentially through the dashboard) would be able to search this catalog to list the available properties they can add to a flavor or image. They will see the virtual CPU topology property in the list and know that it must be an integer. In the dashboard example, when the user adds the property, its key and value will be stored in the service that owns that resource (in nova for flavors, and in glance for images).

openstack-gnocchi

BZ#1252954

> This rebase package addresses the bugs listed in
> https://launchpad.net/gnocchi/+milestone/1.3.0
> #1511656 - gnocchi-metricd traces on empty measures list
> #1496824 - sometimes updating Gnocchi alarms return 400
> #1500646 - data corruption with CEPH and gnocchi-metricd leades to
> delete whole CEPH pool and loose all data
> #1503848 - instance_disk and network_interfaces missing controller
> #1505535 - Intermittent gate failures with py34 + tooz + PG
> #1471169 - MySQL indexer might die with a deadlock
> #1486079 - Delete metric should be async
> #1499372 - Metricd should deal better with corrupted new measures files

```
#1501344 - Evaluation of archive_policy_rules is unspecified
#1504130 - Enhance middleware configuration
#1506628 - Add filter by granularity to measures get
#1499115 - gnocchi-api hangs with api.workers = 2 and CEPH
#1501774 - functional tests fail in the gate with "sudo: .tox/py27-
gate/bin/testr: command not found"
```

openstack-heat

BZ#1303084

Previously, heat would attempt to validate old properties based on the
current property's definitions. Consequently, during director upgrades
where a property definition changed type, the process would fail with a
'TypeError' when heat tried to validate the old property value.
With this fix, heat no longer tries to validate old property values.
As a result, heat can now gracefully handle property schema definitions
changes by only validating new property values.

BZ#1318474

Previously, director used a patch update when updating a cloud, which
reused all the parameters passed at creation. Parameters which were
removed in an update were failing validation. Consequently, updating a
stack with parameters removed, and using a patch update would fail
unless the parameters were explicitly cleared.
With this fix, heat changes the handling of patched updates to ignore
parameters which were not present in the newest template.
As a result, it's now possible to remove top-level parameters and update
a stack using a patch update.

BZ#1303723

Previously, heat would leave the context roles empty when loading the
stored context. When signaling heat used the stored context (trust
scoped token), and if the context did not have any roles, it failed.
Consequently, the process failed with the error 'trustee has no
delegated roles'. This fix addresses this issue by populating roles when
loading the stored context. As a result, loading the auth ref, and
populating the roles from the token will confirm that any RBAC performed
on the context roles will work as expected, and that the stack update
succeeds.

BZ#1303112

Previously, heat changed the name of properties on several neutron
resources; while it used a mechanism to support the old names when
creating them, it failed to validate resources created with a previous
version. Consequently, using Red Hat OpenStack Platform 8 to update a
stack created in version 7 (or previous) using with a neutron port
resource would fail by trying to lookup a 'None' object.
With this fix, when heat updates the resource, it now uses the

> translation mechanism on old properties too. As a result, supporting
> deprecated properties now works as expected with resources created from
> a previous version.

openstack-ironic-python-agent

BZ#1312187

> Sometimes, hard drives were not available in time for a deployment
> ramdisk run. Consequently, the deployment failed if the ramdisk was
> unable to find the required root device. With this update, the "udev
> settle" command is executed before enumerating disks in the ramdisk, and
> the deployment no longer fails due to the missing root device.

openstack-keystone

BZ#1282944

> Identity Service (keystone) used a hard-coded LDAP membership attribute
> when checking if a user was enabled, if the 'enabled emulation' feature
> was being used.
> Consequently, users who were `enabled` could show as `disabled` if an
> unexpected LDAP membership attribute was used.
> With this fix, the 'enabled emulation' membership check now uses the
> configurable LDAP membership attribute that is used for group resources.
> As a result, the 'enabled' status for users is shown correctly when
> different LDAP membership attributes are configured.

BZ#1300395

> This rebase package for Identity Service addresses the following issues:
>
> * Identity Service (keystone) used a hard-coded LDAP membership
> attribute when checking if a user was enabled, if the 'enabled
> emulation' feature was being used. Consequently, users who were
> `enabled` could show as `disabled` if an unexpected LDAP membership
> attribute was used. With this fix, the 'enabled emulation' membership
> check now uses the configurable LDAP membership attribute that is used
> for group resources. As a result, the 'enabled' status for users is
> shown correctly when different LDAP membership attributes are
> configured. (Launchpad bug #1515302, Red Hat BZ#1282944)
>
> * If a user_id just happens to be of 16 character length, the Identity
> service could incorrectly assume that it was handling a UUID value when
> using the Fernet token provider. This would trigger a "Could not find
> user" error in the Identity service logs. This has been corrected to
> properly handle 16 character user IDs. (Launchpad bug #1497461)

BZ#923598

> Previously, the Identity Service (keystone) allowed administrators to
> set a maximum password length limit that was larger than the limit used

by the Passlib python module.
Consequently, if the maximum password length limit was set larger than
the Passlib limit, attempts to set a user password larger than the
Passlib limit would fail with a HTTP 500 response and an uncaught
exception.
With this update, Identity Service now validates that the
'max_password_length' configuration value is less than or equal to the
Passlib maximum password length limit.
As a result, if the Identity Service setting 'max_password_length' is
too large, it will fail to start with a configuration validation error.

openstack-neutron

BZ#1292570

Previously, the 'ip netns list' command returned unexpected ID data in
recent versions of 'iproute2'. Consequently, neutron was unable to parse
namespaces.
This fix addresses this issue by updating the parser used in neutron. As
a result, neutron can now be expected to properly parse namespaces.

BZ#1287736

Prior to this update, the L3 agent failed to respawn keepalived process
if the keepalived parent process died. This was because the child
keepalived process was still running.
Consequently, the L3 agent could not recover from keepalived parent
process death, breaking the HA router served by the process.
With this update, the L3 agent is made aware of the child keepalived
process, and now cleans up it as well before respawning keepalived.
As a result, the L3 agent is now able to recover HA routers when the
keepalived process dies.

BZ#1290562

Red Hat OpenStack Platform 8 introduced a new RBAC feature that allows
you to share neutron networks with a specific list of tenants, instead
of globally. As part of the feature, the default policy.json file for
neutron started triggering I/O, consuming database fetches for every
port fetch in attempt to allow the owner of a network to list all ports
that belong to his network, even if they were created by other tenants.
Consequently, the list operation for ports triggered multiple unneeded
database fetches, which drastically affected performance of the
operation.
This update addresses this issue by running the I/O operations only when
they are actually needed, for example, when the port to be validated by
the policy engine does not belong to the tenant that invokes the list
operation. As a result, list operations for ports will scale normally
again.

BZ#1222775

Prior to this update, the fix for BZ#1215177 added the
'garp_master_repeat 5' and 'garp_master_refresh 10' options to
Keepalived configuration.
Consequently however, Keepalived continuously spammed the network with
Gratuitous ARP (GARP) broadcasts; in addition, instances would lose
their IPv6 default gateway settings. As a result of these issues, the
IPv6 router stopped working with VRRP.
This update addresses these issues by dropping the 'repeat' and
'refresh' Keepalived options. This fixes the IPv6 bug but re-introduces
the bug described in BZ#1215177.
To resolve this, use the 'delay' option instead. As a result, Keepalived
sends a GARP when it transitions to 'MASTER', and then waits a number of
seconds (determined by the delay option), and sends another GARP. Use an
aggressive 'delay' setting to make sure that when the node boots and the
L3/L2 agents start, there is enough time for the L2 agent to wire the
ports.

BZ#1283623

Prior to this update, a change to the Open vSwitch agent introduced a
bug in how the agent handles the segmentation ID value for flat
networking during agent startup.
Consequently, the agent failed to restart when serving a flat network.
With this update, the agent code was fixed to handle segmentation
properly for flat networking. As a result, the agent is successfully
restarted when serving a flat network.

BZ#1295690

Previously, a router that was neither an HA nor a DVR router could not
be converted into an HA router. Instead, it was necessary to create a
new router and reconnect all the resources (interfaces, networks etc.)
from the old router to the new one. This update adds the ability to
convert a legacy router into an HA or non-HA router in a few simple
commands:

```
# neutron router-update ROUTER --admin-state-up=False
# neutron router-update ROUTER --ha=True/False
# neutron router-upgrade ROUTER --admin-state-up=True
```

Replace ROUTER with the ID or name of the router to convert.

BZ#1177611

A known issue has been identified for interactions between High
Availability (VRRP) routers and L2 Population. Currently, when
connecting a HA router to a subnet, HA routers use a distributed port by
design. Each router has the same port details on each node that it's
scheduled on, and only the master router has IPs configured on that
port; all the slaves have the port without any IPs configured.
Consequently, L2Population uses the stale information to advise that the
router is present on the node (which it states in the port binding
information for that port).
As a result, each node that has a port on that logical network has a

tunnel created only to the node where the port is presumably bound. In addition, a forwarding entry is set so that any traffic to that port is sent through the created tunnel.
However, this action may not succeed as there is not guarantee that the master router is on the node specified in the port binding. Furthermore, in the event that the master router is in fact on the node, a failover event would cause it to migrate to another node and result in a loss of connectivity with the router.

BZ#1300308

Previously, the neutron-server service would sometimes erroneously require a new RPC entrypoint version from the L2 agents that listened for security group updates.
Consequently, the RHEL OpenStack Platform 7 neutron L2 agents could not handle certain security group update notifications sent by Red Hat OpenStack Platform 8 neutron-server services, causing certain security group updates to not be propagated to the data plane.
This update addresses this issue by ending the requirement of the new RPC endpoint version from agents, as this will assist the rolling upgrade scenario between RHEL OpenStack Platform 7 and Red Hat OpenStack Platform 8.
As a result, RHEL OpenStack Platform 7 neutron L2 agents will now correctly handle security group update notifications sent by the Red Hat OpenStack Platform 8 neutron-server services.

BZ#1293381

Prior to this update, when the last HA router of a tenant was deleted, the HA network belonging to the tenant was not removed. This happened in certain scenarios, such as the 'router delete' API call, which raised an exception since the router had been deleted. That scenario was possible due to a race condition between HA router 'create' and 'delete' operations. As a result of this issue, HA network tenants were not deleted.
This update resolves the race condition, and now catches the exceptions 'ObjectDeletedError' and 'NetworkInUse' when a user deletes the last HA router, and also moves the HA network deleting procedure under the 'ha_network exist' check block. In addition, the fix checks whether or not HA routers are present, and deletes the HA network when the last HA router is deleted.

BZ#1255037

Neutron ports created when neutron-openvswitch-agent is down are in status "DOWN, binding:vif_type=binding_failed", which is expected. Nevertheless, prior to this update, there was no way to recover those ports even if neutron-openvswitch-agent was back online. Now, the function "_bind_port_if_needed" binds at least once when the port's binding status passed in is already in "binding_failed". As a result, ports can now recover from a failed binding status by repeated binding attempts triggered when neutron-openvswitch-agent comes back online.

BZ#1284739

Prior to this update, the status of a floating IP address was not set
when the floating IP address was realized by an HA router. Consequently,
'neutron floatingip-show <floating_ip>' would not output an updated
status.
With this update, a floating IP address status is updated when realized
by HA routers, and when the L3 agent configures a router.
As a result, the status field for floating IP addresses realized by HA
routers are now updated to 'ACTIVE' when the floating IP is configured
by the L3 agent.

openstack-nova

BZ#978365

The ability of the libvirt driver to set the admin password has been
added. To use this feature, run the following command: "nova root-
password [server]".

BZ#1298825

Previously, selecting an odd number of vCPUs would cause the assignment
of one core and one thread in the guest instance per CPU, which would
impact performance.
The update addresses this issue by correctly assigning pairs of threads
and one independent thread per CPU, when an odd number of vCPUs is
assigned.

BZ#1301914

Previously, when a source compute node is back up after a migration,
instances that have been successfully evacuated from it when the node
was down were not deleted. A result of having the non-deleted instances
makes it impossible to evacuate them.

With this update, the successful migration status when evacuating an
instance is now verified for knowing which instance to delete when a
compute node is back up and running again. As a result, instances can be
evacuated from one host to another, regardless of their previous
locations.

BZ#1315394

This package rebases Compute (nova) to version 12.0.2, and includes a
number of updates:
- Propagate qemu-img errors to compute manager
- Fix evacuate support with Nova cells v1
- libvirt: set libvirt.sysinfo_serial='none' for virt driver tests
- XenAPI: Workaround for 6.5 iSCSI bug
- Change warn to debug logs when migration context is missing
- Imported Translations from Zanata

- libvirt: Fix/implement revert-resize for RBD-backed images
- Ensure Glance image 'size' attribute is 0, not 'None'
- Add retry logic for detaching device using LibVirt
- Spread allocations of fixed ips
- Apply scheduler limits to Exact* filters
- Replace eventlet-based raw socket client with requests
- VMware: Handle image size correctly for OVA and streamOptimized images
- XenAPI: Cope with more Cinder backends
- ports and networks gather should validate existance
- Disable IPv6 on bridge devices
- Validate translations
- Fix instance not destroyed after successful evacuation

BZ#1293607

With this update, the 'openstack-nova' packages have been rebased to upstream version 12.0.1.

Some of the highlights addressed by this rebase are as follows:

- Treat sphinx warnings as errors when building release notes
- Fix warning in 12.0.1-cve-bugs-7b04b2e34a3e9a70.yaml release note
- Fix backing file detection in libvirt live snapshot
- Add security fixes to the release notes for 12.0.1
- Fix format conversion in libvirt snapshot
- Fix format detection in libvirt snapshot
- VMware: specify chunk size when reading image data
- Revert "Fixes Python 3 str issue in ConfigDrive creation"
- Do not load deleted instances
- Make scheduler_hints schema allow list of id
- Add -constraints sections for CI jobs
- Remove the TestRemoteObject class
- Update from global requirements
- VMware: fix bug for config drive when inventory folder is used
- Omnibus stable fix for upstream requirements breaks
- Refresh stale volume BDMs in terminate_connection
- Fix metadata service security-groups when using Neutron
- Add "vnc" option group for sample nova.conf file
- Scheduler: honor the glance metadata for hypervisor details
- reno: document fixes for service state reporting issues
- servicegroup: stop zombie service due to exception
- Import Translations from Zanata
- xen: mask passwords in volume connection_data dict
- Fix is_volume_backed_instance() for unset image_ref
- Split up test_is_volume_backed_instance() into five functions
- Handle DB failures in servicegroup DB driver
- Fixes Python 3 str issue in ConfigDrive creation
- Fix Nova's indirection fixture override
- Updated from global requirements
- Add first reno-based release note
- Add "unreleased" release notes page
- Add reno for release notes management
- The test_schedule_to_all_nodes test is currently broken and has been blacklisted.
- libvirt:on snapshot delete, use qemu-img to blockRebase if VM is stopped

- Fix attitute error when cloning raw images in Ceph
- Exclude all BDM checks for cells
- Image meta: treat legacy vmware adapter type values

openstack-packstack

BZ#1301366

Previously, Packstack did not enable the VPNaaS tab in the Dashboard
even if the CONFIG_NEUTRON_VPNAAS parameter was set to 'y'. As a result,
the tab for VPNaaS was not shown on the Dashboard.

With this update, a check to see if VPNaaS is enabled has been set up.
This check then enables the Dashboard tab in the Puppet manifest. As a
result, the VPNaaS tab is now shown on the Dashboard when the service is
configured in Packstack.

BZ#1297712

Previously, Packstack edited the /etc/lvm/lvm.conf file to set specific
parameters for snapshot autoextend. However, the regexp used only
allowed black spaces instead of the tabs as currently used in the file.
As a result, some lines were added at the end of the file, breaking its
format.

With this update, the regexp is updated in Packstack to set the
parameters properly. As a result, there are no error messages when
running LVM commands.

openstack-puppet-modules

BZ#1289180

Previously, although the haproxy is configured to allow a value of 10000
for the 'maxconn' parameter for all proxies together, there is a default
'maxconn' value of 2000 for each proxy individually. If the specific
proxy used for MySQL reached the limit of 2000, it dropped all further
connections to the database and the client would not retry, which caused
API timeout and subsequent commands to fail.

With this update, the default value for 'maxconn' parameter has been
increased to work better for production environments, As a result, the
database connections are far less likely to time out.

BZ#1280523

Previously, Facter 2 did not have netmask6 and netmask6_<ifce> facts. As
a result, IPv6 was not supported.

With this update, the relevant custom facts have been added to support checks on IPv6 interfaces, resulting in the IPv6 interfaces are now supported.

BZ#1243611

Previously, there was no default time out parameter, resulting in some stages of Ceph cluster set-up that look longer than the default 5 minutes (300 seconds).

With this update, a time out parameter is added for relevant operations. The default time out parameter value is set at 600 seconds. You can modify the default value, if necessary. As a result, the installation is more resilient, especially when some of the Ceph setup operations take longer than average.

openstack-sahara

BZ#1189502

With this update, configuration settings now exist to set timeouts, after which clusters which have failed to reach the 'Active' state will be automatically deleted.

BZ#1189517

When creating a job template intended for re-use, you can now register a variable for datasource URLs with OpenStack Data Processing (sahara). Doing so allows you to easily change input and output paths per run, rather than an actual URL (which would require revising the template, or manually revising the URL per run between jobs).

This makes it easier to reuse job templates when data source jobs are mutable between runs, as is true for most real-world cases.

BZ#1299982

With this update, the integration for CDH 5.4 with Sahara is now complete and hence, the default-enabled option for the plugin version, CDH 5.3 is now removed.

BZ#1233159

Previously, the tenant context information was not available to the periodic task responsible for cleaning up stale clusters.

With this update, temporary trusts are established between the tenant and admin, allowing the periodic job to use this trust to delete stale clusters.

openstack-selinux

BZ#1281547

Previously, httpd was not allowed to search through directories having the "nova_t" label. Consequently, nova-novncproxy failed to deploy an HA overcloud. This update allows httpd to search through such directories, which enables nova-novncproxy to run successfully.

BZ#1284268

Previously, Openvswitch was trying to create a tun socket, but SELinux prevented that. This update allows Openvswitch to create a tun socket, and as a result, Openvswitch now runs without failures.

BZ#1310383

Previously, SELinux blocked ovsdb-server from running, resulting in simple networking operations to fail.

With this update, Open vSwitch is allowed to connect to its own port. As a result, ovsdb-server now runs without issues and the networking operations are completed successfully.

BZ#1284133

Previously, SELinux prevented redis from connecting to its own port, resulting in redis failing at restart.

With this update, redis has the permission to connect to the 'redis' labeled port. As a result, redis runs properly and resource restart is successful.

BZ#1281588

Prior to this update, SELinux prevented nova from uploading the public key to the overcloud. A new rule has now been added to allow nova to upload the key.

BZ#1306525

Previously, when nova was trying to retrieve a list of glance images, SELinux prevented that, and nova failed with an "Unexpected API Error". This update allows nova to communicate with glance. As a result, nova can now list glance images.

BZ#1283674

Prior to this update, SELinix prevented dhclient, vnc, and redis from working. New rules have now been added to allow these software tools to run successfully.

openvswitch

BZ#1266050

> The Open vSwitch (openvswitch) package is now re-based to upstream
> version 2.4.0.

python-cinderclient

BZ#1214230

> With this update, a new feature adds pagination support for the Block
> Storage 'snapshots-list' and 'backups-list' commands. You can now limit,
> marker and sort parameters to control the number of returned results,
> starting element and their order.
>
> Retrieving a limited number of results instead of the entire data set
> can be extremely useful on the large deployments with thousands of
> snapshots and backups.

python-django-horizon

BZ#1167563

> The 'Launch Instance' workflow has been redesigned and re-implemented to
> be more responsive with this update.
>
> 1. To enable this update, add the following values in your
> /etc/openstack-dashboard/local_settings file:
>
> LAUNCH_INSTANCE_LEGACY_ENABLED = False
> LAUNCH_INSTANCE_NG_ENABLED = True
>
> 2. Restart 'httpd':
> # systemctl restart httpd

BZ#1100542

> OpenStack dashboard tables summarize information about a large number of
> entities. This update adds a table enhancement that enables this
> information to be displayed within the table as a slide-down "drawer"
> that is activated when you click on a toggle switch within a row. The
> drawer appears as an additional row (with configurable height) and
> contains additional information about the entity in the row above it
> (e.g. additional entity details, metrics, graphs, etc.). Multiple
> drawers may be opened at one time.

BZ#1166963

> This update replaces the network topology with curvature based graph as
> the previous UI did not work well with larger number of nodes or
> networks.

The new network topology map can handle more nodes, looks stylish and the node layout can be re-organized.

BZ#1042947

This update adds support for volume migrations of the Block Storage (cinder) service. These are done in the 'Volumes' panel of the OpenStack dashboard (Project-> Compute -> Volumes and in Admin-> System Panel-> Volumes). You can perform this action on the 'Volumes' row in the table. The final patch in this series resolved the command action itself; it had previously errored out due to incorrect parameters, and parameter count issues.

BZ#1305905

The python-django-horizon packages have been upgraded to upstream version 8.0.1, which provides a number of bug fixes and enhancements over the previous version. Notably, this version contains localization updates, includes Italian localization, fixes job_binaries deletion, and adds support for accepting IPv6 in the VIP address for an LB pool.

BZ#1279812

With this release, panels are configurable. You can add or remove panels by using configuration snippets.

For example, to remove the "Resource panel":

* Place a file in '/usr/share/openstack-dashboard/openstack_dashboard/local/enabled'.
* Name that file '_99_disable_metering_dashboard.py'.
* Copy the following content into the file:

```
# The slug of the panel to be added to HORIZON_CONFIG. Required.
PANEL = 'metering'
# The slug of the dashboard the PANEL associated with. Required.
PANEL_DASHBOARD = 'admin'
# The slug of the panel group the PANEL is associated with.
PANEL_GROUP = 'admin'
REMOVE_PANEL = True
```

* Restart the Dashboard httpd service:
# systemctl restart httpd

For more information, see the Pluggable Dashboard Settings

BZ#1300735

With this release, the 'Metering' panel in Dashboard (horizon) has been disabled due to performance issues.

BZ#1297757

> Previously, no timeout was specified in horizon's systemd snippet for
> httpd, so the standard one-minute timeout was used when waiting for
> httpd to fully start up. In some cases, however, especially when running
> in a virtualized or a very loaded environment, the startup takes longer.
> Consequently, a failure from systemd sometimes occurred even if httpd
> was already running. With this update, the timeout has been set to two
> minutes, which resolves the problem.

python-glance-store

BZ#1284845

> Previously, when Object Storage service was used as a backend storage
> for Image service, image data was stored in Object Storage service as
> multiple 'chunks' of data. When using the Image service APIv2, there
> were circumstances in which the upload operations would fail if the
> client sent a final zero-sized 'chunk' to the server. The failure
> involved a race condition between the operation to store a zero-sized
> 'chunk' and a cleanup delete of that 'chunk'. As a result, intermittent
> failure occurred while storing Image service images in Object Storage
> service.
>
> With this update, the cleanup delete operations are retried rather than
> failing them as well as the primary upload image task. As a result,
> Image service APIv2 handles this rare circumstance gracefully, so that
> the image upload does not fail.

BZ#1229634

> Previously, there was no secure way to remotely access S3 backend in a
> private network.
>
> With this update, a new feature allows Image service S3 driver to
> connect a S3 backend from a different network in a secure way through
> the HTTP proxy.

python-glanceclient

BZ#1314069

> Previously, the Image service client could be configured to only allow
> uploading images in certain formats (for example, raw, ami, iso) to the
> Image service server. The client also allowed download of an image from
> the server only if it was in one of these formats. As a result of this
> restriction, users could no longer download images in other formats that
> had been previously uploaded.
>
> With this update, as the Image service server already validates image
> formats at the time they are imported, there is no need for the Image
> service client to verify image format when it is downloaded. As a

result, the image format validation when an image is downloaded is now skipped, allowing the consumption of images in legitimate formats even if the client-side support for upload of images in those formats is no longer configured.

python-heatclient

BZ#1234108

Previously, the output of the "heat resource-list --nested-depth ..." command contained a column called "parent_resource"; however, the output did not include the information required to run a subsequent "heat resource-show ..." command. With this update, the output of the "heat resource-list --nested-depth ..." command includes a column called "stack_name", which provides the values to use in a "heat resource-show [stack_name] [resource_name]" call.

python-networking-odl

BZ#1266156

The OpenDaylight OpenStack neutron driver has been split from the neutron project and moved to a new package, python-networking-odl. Operators still have the driver available for use as part of their Red Hat OpenStack Platform installations.

python-neutronclient

BZ#1291739

The 'neutron router-gateway-set' command now supports the '--fixed-ip' option, which allows you to configure the fixed IP address and subnet that the router will use in the external network. This IP address is used by the OpenStack Networking service (openstack-neutron) to connect interfaces on the software level to connect the tenant networks to the external network.

python-openstackclient

BZ#1303038

With this release, the python-openstackclient package is now re-based to upstream version 1.7.2. This applies several fixes and enhancements, which include improved exception handling for 'find_resource'.

python-oslo-messaging

BZ#1302391

Oslo Messaging used the "shuffle" strategy to select a RabbitMQ host from the list of RabbitMQ servers. When a node of the cluster running RabbitMQ was restarted, each OpenStack service connected to this server reconnected to a new RabbitMQ server. Unfortunately, this strategy does not handle dead RabbitMQ servers correctly; it can try to connect to the same dead server multiple times in a row. The strategy also leads to increased reconnection time, and sometimes it may lead to RPC operations timing out because no guarantee is provided on how long the reconnection process will take.

With this update, Oslo Messaging uses the "round-robin" strategy to select a RabbitMQ host. This strategy provides the least achievable reconnection time and avoids RPC timeout when a node is restarted. It also guarantees that if K of N RabbitMQ hosts are alive, it will take at most N - K + 1 attempts to successfully reconnect to the RabbitMQ cluster.

BZ#1312912

When the RabbitMQ service fails to deliver an AMQP message from one OpenStack service to another, it reconnects and retries delivery. The "rabbit_retry_backoff" option, whose default is 2 seconds, is supposed to control the pace of retries; however, retries were previously done every second irrespective of the configured value of this option. The consequence of this problem was excessive retries, for example, when an endpoint was not available. This problem has now been fixed, and the "rabbit_retry_backoff" option, as explicitly configured or with the default value of two seconds, properly controls message delivery retries.

python-oslo-middleware

BZ#1313875

With this release, oslo.middleware now supports SSL/TLS, which in turn allows OpenStack services to listen to HTTPS traffic and encrypt exchanges. In previous releases, OpenStack services could only listen to HTTP, and all exchanges were done in cleartext.

python-oslo-service

BZ#1288528

A race condition in the SIGTERM and SIGINT signal handlers made it possible for worker processes to ignore incoming SIGTERM signals. When two SIGTERM signals were received "quickly" in child processes of OpenStack services, some worker processes could fail to handle incoming SIGTERM signals; as a result, those processes would remain active. Whenever this occurred, the following AssertionError exception message appeared in logs:

```
Cannot switch to MAINLOOP from MAINLOOP
```

> This release includes an oslo.service that fixes the race condition,
> thereby ensuring that SIGTERM signals are handled correctly.

sahara-image-elements

BZ#1286276

> In some base image contexts, iptables was not initialized prior to save.
> This cause 'iptables save' in the 'disable-firewall' element to fail.
> This release adds the non-destructive command 'iptables -L', which
> successfully initializes iptables in all contexts, thereby ensuring a
> successful image generation.

BZ#1286856

> In the Liberty release, the OpenStack versioning scheme is now based on
> the major release number (previously, it was based on year). This update
> adds an epoch to the current sahara-image-elements package to ensure
> that it upgrades the older version.

4.2. RHEA-2016:0604 - RED HAT OPENSTACK PLATFORM 8 DIRECTOR ENHANCEMENT ADVISORY

The bugs contained in this section are addressed by advisory RHEA-2016:0604. Further information about this advisory is available at https://access.redhat.com/errata/RHEA-2016:0604.html.

instack-undercloud

BZ#1212158

> This updates provides OpenStack notifications. Previously there were
> external consumers of OpenStack notifications that could not interface
> with director-deployed cloud because notifications were not enabled. Now
> director enables notifications for external consumers.

BZ#1223257

> A misconfiguration of Ceilometer on the Undercloud caused hardware
> meters to not work correctly. This fix provides a vaild default
> Ceilometer configuration. Now Ceilometer hardware meters work as
> expected.

BZ#1296295

> Running "openstack undercloud install" attempted to delete and recreate
> the Undercloud's neutron subnet even if the subnet required no changes.
> If an Overcloud was already deployed, the subnet delete attempt failed
> since the subnet contained allocated ports. This caused the "openstack
> undercloud install" command to fail. This fix changes this behavior to

only attempt to delete and recreate the subnet if the "openstack undercloud install" command has a configuration change to apply to the subnet. If an Overcloud is already deployed, the same error message still occurs since the director cannot delete the subnet. This is expected behavior though since we do not recommend change the subnet's configuration with an Overcloud already deployed. However, in cases with no subnet configuration changes, the "openstack undercloud install" command no longer fails with this error message.

BZ#1298189

The Puppet manifest that installs the Undercloud referred to the wrong resource name to create the keystone domain for Heat. The undercloud install failed with an error such as:

puppet apply exited with exit code 1

This fix updates the Puppet manifest was changed to use the correct name of the resource. The Undercloud installation now finishes without an error.

BZ#1315546

When LANG was set to ja_JP.UTF-8, the output of the date command in "dib-run-parts" contained Japanese characters, which caused a unicode error in "_run_live_command()". The Undercloud installation failed with the following error:

UnicodeDecodeError: 'ascii' codec can't decode byte 0xe5 in position 18: ordinal not in range(128)
Command 'instack-install-undercloud' returned non-zero exit status 1

This fix decodes strings using the utf-8 character encoding during the Undercloud installation. Now the Undercloud installation completes successfully when LANG is set to ja_JP.UTF-8.

BZ#1312889

This update removes the Tuskar API service from Red Hat OpenStack Platform director 8. Tuskar was installed and configured on the Undercloud, including an endpoint existing in the Keystone service catalog. The RPM is no longer installed, the service is not configured, and the endpoint is not created in the service catalog.

openstack-ironic-inspector

BZ#1282580

The director includes new functionality to allow automatic profile matching. Users can specify automatic matching between nodes and deployment roles based on data available from the introspection step.

> Users now use ironic-inspector introspection rules and new python-
> tripleoclient commands to assign profiles to nodes.

BZ#1270117

> Previously, periodic iptables calls made by Ironic Inspector did not
> contain the -w option, which instructs iptables to wait for the xtables
> lock. As a consequence, periodic iptables updates occasionally failed.
> This update adds the -w option to the iptables calls, which prevents the
> periodic iptables updates from failing.

openstack-ironic-python-agent

BZ#1283650

> Log processing in the introspection ramdisk did not take into account
> non-Latin characters in logs. Consequently, the "logs" collector failed
> during introspection. With this update, log processing has been fixed to
> properly handle any encoding.

BZ#1314642

> The director uses a new ramdisk for inspection and deployment. This
> ramdisk included a new algorithm to pick the default root device for
> users not using root device hints. However, possible root device changes
> occurred on redployment, leading to failures. This fix reverts the
> ramdisk device logic to be the same as OpenStack Platform director 7.
> Note that this does not mean that the default root device is the same,
> as device names are not reliable. Also this behavior will change again
> in a future releases. Make sure to use root device hints if you nodes
> use multiple hard drives.

openstack-tripleo-heat-templates

BZ#1295830

> Pacemaker used a 100s timeout for service resources. However, a systemd
> timeout requires an additional timeout period after the initial timeout
> to accommodate for a SIGTERM and then a SIGKILL. This fix increases the
> Pacemaker timeout to 200s to accommodate two full systemd timeout
> periods. Now the timeout period is enough for systemd to perform a
> SIGTERM and then a SIGKILL.

BZ#1311005

> The notify=true parameter was previously missing from the RabbitMQ
> Pacemaker resource. Consequently, RabbitMQ instances were unable to
> rejoin the RabbitMQ cluster. This update adds support for notify=true to
> the pacemaker resource agent for RabbitMQ, and adds notify=true to
> OpenStack director. As a result, RabbitMQ instances are now able to
> rejoin the RabbitMQ cluster.

■

BZ#1283632

The 'ceilometer' user lacked a role needed for some functionality, which causes some Ceilometer meters to function incorrectly. This fix adds the necessary role to the 'ceilometer' user. Now all ceilometer meters work correctly.

BZ#1299227

Prior to this update, the swift_device and swift_proxy_memcache URIs used for the swift ringbuilder and the swift proxy memcache server respectively were not properly formatted for IPv6 addresses, lacking the expected '[]' delimiting the IPv6 address. As a consequence, when deploying with IPv6 enabled for the overcloud, the deploy failed with "Error: Parameter name failed on Ring_object_device ...". Now, when IPv6 is enabled, the IP addresses used as part of the swift_device and swift_proxy_memcache URIs are correctly delimited with '[]'. As a result, deploying with IPv6 no longer fails on incorrect formatting for swift_device or swift_proxy_memcache.

BZ#1238807

This enhancement enables the distribution of per-node hieradata, matching the nodes from their UUID (as reported by 'dmidecode'). This allows you to scale CephStorage across nodes equipped with a different number/type of disks.
As a result, CephStorage nodes can now be configured with non-homogeneous disk topologies. This is done by provisioning a different configuration hash for the ceph::profile::params::osds parameter.

BZ#1242396

Previously, the os-collect-config utility only printed Puppet logs after Puppet had finished running. As a consequence, Puppet logs were not available for Puppet runs that were in progress. With this update, logs for Puppet runs are available even when a Puppet run is in progress. They can be found in the /var/run/heat-config/deployed/ directory.

BZ#1266104

This update adds neutron QoS (Quality of Service) extensions to provide better control over tenant networking qualities and limits. Overclouds are now deployed with Neutron QoS extension enabled.

BZ#1320454

Stricter validation in Red Hat OpenStack Platform 8's Orchestration service (heat) caused the Overcloud stack update to fail from an upgraded Undercloud with the following error:

ERROR heat.engine.resource ResourceFailure: resources.Compute:

"u'1:1000'" is not a list.

This fix which properly formats the NeutronVniRanges parameter to include the required '[]' has been backported to the OpenStack Platform 7 openstack-tripleo-heat-templates package and should be available as of openstack-tripleo-heat-templates-kilo-0.8.14-5.el7ost.noarch. Now stack updates of the Overcloud stack will not fail with the error when using an upgraded Undercloud to manage an existing Overcloud while using the version 7 templates (which are located at /usr/share/openstack-tripleo-heat-templates/kilo).

BZ#1279615

This update allows enabling of the Neutron L2 population feature. This helps reduce the amount of broadcast traffic in Tenant networks. Set the NeutronEnableL2Pop parameter in an environment file's 'default_parameters' section to enable Neutron L2 population.

BZ#1225163

The Director now properly enabled notifications for external consumers.

BZ#1259003

The domain name for overcloud nodes defaulted to 'localdomain'. For example: 'overcloud-compute-0.localdomain'. This enhancement provides a parameter (CloudDomain) to customize the domain name. Create an environment file with the CloudDomain parameter included in the 'parameter_defaults" section. If no domain name is defined, the Heat templates default to 'localdomain'.

BZ#1273303

The Director now supports the OpenStack Networking 'enable_isolated_metadata' option. This option allows access to instance metadata on VMs on external routers or on isolated networks.

BZ#1308422

Previously, '/v2.0' was missing from the end of the URL specified in the admin_auth_url setting in the [neutron] section of /etc/nova/nova.conf. This would prevent Nova from being able to boot instances because it could not connect to the Keystone catalog to query for the Neutron service endpoint to create and bind the port for instances. Now, '/v2.0' is correctly added to the end of the URL specified in the admin_auth_url setting, allowing instances to be started successfully after deploying an overcloud with the director.

BZ#1298247

The Director now supports new parameters that control whether to disable or enable the following OpenStack Networking services:

```
* dhcp_agent
* l3_agent
* ovs_agent
* metadata_agent

This enhancement allows the deployment of Neutron plug-ins that replace
any of these services. To disable all of these services, use the
following parameters in your environment file:

  NeutronEnableDHCPAgent: false
  NeutronEnableL3Agent: false
  NeutronEnableMetadataAgent: false
  NeutronEnableOVSAgent: false
```

BZ#1266219

```
The Director can now deploy the Block Storage service with a Dell
EqualLogic or Dell Storage Center appliance as a back end. For more
information, see:

https://access.redhat.com/documentation/en/red-hat-openstack-
platform/version-8/dell-equallogic-back-end-guide/
https://access.redhat.com/documentation/en/red-hat-openstack-
platform/8/dell-storage-center-back-end-guide/dell-storage-center-back-
end-guide
```

os-cloud-config

BZ#1288475

```
A bug in the Identity service's endpoint registration code failed to
mark the Telemetry service as SSL-enabled. This prevented the Telemetry
service endpoint from being registered as HTTPS. This update fixes the
bug: the Identity service now correctly registers Telemetry, and
Telemetry traffic is now encrypted as expected.
```

BZ#1319878

```
When using Linux kernel mode for bridges and bonds (as opposed to Open
vSwitch), the physical device was not detected for the VLAN interfaces.
This, in turn, prevented the VLAN interfaces from working correctly.

With this release, the os-net-config utility automatically detects the
physical interface for a VLAN as long as the VLAN is a member of the
physical bridge (that is, the VLAN must be in the 'members:' section of
the bridge). As such, VLAN interfaces now work properly with both OVS
bridges and Linux kernel bridges.
```

BZ#1316730

```
In previous releases, when VLAN interfaces were placed directly on a
```

Linux kernel bond with no bridge, it was possible for the VLANs to start
before the bond. When this occurred, the VLANs failed to start. With
this release, the os-net-config utility now starts the physical network
(namely, bridges first, then bonds and interfaces) before VLANs. This
ensures that the VLANs have the interfaces necessary to start properly.

python-rdomanager-oscplugin

BZ#1271250

In previous releases, a bug made it possible for failed nodes to be
marked as available. Whenever this occurred, deployments failed because
nodes were not in a proper state. This update backports an upstream
patch to fix the bug.

python-tripleoclient

BZ#1288544

Previously, bulk introspection only printed on-screen errors, but never
returned a failure status code. This prevented introspection failures
from being detected. This update changes the status code of errors to
non-zero, which ensures that failed introspections can now be detected
through their status codes.

BZ#1261920

Previously, bulk introspection operated on nodes currently in
maintenance mode. This could cause introspection to fail, or even break
node maintenance (depending on the reason for node maintenance). With
this release, bulk introspection now ignores nodes in maintenance mode.

BZ#1246589

In older deployments using the python-rdomanager-oscplugin (not the
python-tripleoclient) for Overcloud deployment, the
dhcp_agents_per_network parameter for neutron was set to a minimum of 3,
even in the case of a non-HA single controller deployment. This meant
the dhcp_agents_per_network was set to 3 when deploying with only 1
Controller. This fix takes into account the single Controller case. The
director sets at most 3 dhcp_agents_per_network and never more than the
number of Controllers. Now if you deploy in HA with 3 or more controller
nodes, the dhcp_agents_per_network configuration parameter in
neutron.conf on those Controller nodes will be set to '3'. Alternatively
if you deploy in non-HA with only 1 Controller, this same
dhcp_agents_per_network parameter will be set to '1'.

rhel-osp-director

BZ#1293979

Updating packages on the Undercloud left the Undercloud in an indeterminate state. This meant some Undercloud services were disabled after the package update and could not start again. As a workaround, run 'openstack undercloud install' to reconfigure all Undercloud services. After the command complete, the Undercloud services operate normally.

BZ#1234601

The Ramdisk and Kernel images booted without specifying a particular interface. This meant the system booted from any network adapter, which caused problems when more than one interface was on the Provisioning network. In those cases it was necessary to specify which interface the system should use to boot. The interface specified should correspond to the interface which carried the MAC address from the instackenv.json file.

As a workaround, copy and paste the following block of text as the root user into the director's terminal.This creates a systemd startup script sets these parameters on every boot.

The script contains a sed command which includes "net0/mac". This sets the director to use the first Ethernet interface. Change this to "net1/mac" to use the second interface, and so on.

```
#######################################
cat << EOF > /usr/bin/bootif-fix
#!/usr/bin/env bash

while true;
        do find /httpboot/ -type f ! -iname "kernel" ! -iname "ramdisk"
! -iname "*.kernel" ! -iname "*.ramdisk" -exec sed -i 's|{mac|
{net0/mac|g' {} +;
done
EOF

chmod a+x /usr/bin/bootif-fix

cat << EOF > /usr/lib/systemd/system/bootif-fix.service
[Unit]
Description=Automated fix for incorrect iPXE BOOFIF

[Service]
Type=simple
ExecStart=/usr/bin/bootif-fix

[Install]
WantedBy=multi-user.target
EOF

systemctl daemon-reload
systemctl enable bootif-fix
systemctl start bootif-fix

#########################################
```

> The bootif-fix script runs on every boot. This enables booting from a specified NIC when more than one NIC is on the Provisioning network. To disable the service and return to the previous behavior, run "systemctl disable bootif-fix" and reboot.

BZ#1249601

> OpenStack Bare Metal (ironic) now supports deploying nodes in UEFI mode. This is due to requests from customers with servers that only support UEFI boot.

BZ#1236372

> A misconfiguration of the health check for Nova EC2 API caused HAProxy to believe the API was down. This meant the API was unreachable through HAProxy. This fix corrects the health check to query the API service state correctly. Now the Nova EC2 API is reachable through HAProxy.

BZ#1265180

> The director requires the 'baremetal' flavor, even if unused. Without this flavor, the deployment fails with an error. Now the Undercloud installation automatically creates the 'baremetal' flavor. With the flavor in place, the director does not report the error.

BZ#1318583

> Previously, the os_tenant_name variable in the Ceilometer configuration was incorrectly set to the 'admin' tenant instead of the 'service' tenant. This caused the ceilometer-central-agent to fail with the error "ERROR ceilometer.agent.manager Skipping tenant, keystone issue: User 739a3abf8504498e91044d6d2a6830b1 is unauthorized for tenant d097e6c45c494c2cbef4071c2c273a58". Now, Ceilometer is correctly configured to use the 'service' tenant.

BZ#1315467

> Previously, after upgrading the undercloud, there was a missing restart of the openstack-nova-api service, which would cause upgrades of the overcloud to fail due to a timeout that would report the error "ERROR: Timed out waiting for a reply to message ID 84a44ca3ed724eda991ba689cc364852". Now, the openstack-nova-api service is correctly restarted as part of the undercloud upgrade process, allowing the overcloud upgrade process to proceed without encountering this timeout issue.

4.3. RHBA-2016:1063 - OPENSTACK-NEUTRON BUG FIX ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2016:1063. Further information about this advisory is available at https://access.redhat.com/errata/RHBA-2016:1063.html.

4.3.1. openstack-neutron

BZ#1286302

> Previously, using 'neutron-netns-cleanup' when manually taking down a
> node from an HA cluster would not properly clean up processes in the
> neutron L3-HA routers. Consequently, when the node was connected again
> to the cluster, and services were re-created, the processes would not
> properly respawn with the right connectivity. As a result, even if the
> processes were alive, they were disconnected; this sometimes led to a
> situation where no L3-HA router was able to take the 'ACTIVE' role.
> With this update, the 'neutron-netns-cleanup' scripts and related OCF
> resources have been fixed to kill the relevant keepalived processes and
> child processes.
> As a result, nodes can be taken off the cluster and back, and the
> resources will be properly cleaned up when taken off the cluster, and
> restored when taken back.

BZ#1325806

> With this update, OpenStack Networking has been rebased to version
> 7.0.4.
>
> This update introduces the following enhancements:
>
> * Add an option for nova endpoint type
> * Update devstack plugin for dependent packages
> * De-duplicate conntrack deletions before running them
> * Unmarshall portinfo on update_fdb_entries calls
> * Avoid DuplicateOptError in functional tests
> * Retry port create/update on duplicate db records
> * Catch PortNotFound after HA router race condition
> * Documenting network_device_mtu in agents config files
> * Make all tox targets constrained
> * Filter HA routers without HA interface and state
> * Correct return values for bridge sysctl calls
> * Add tests for RPC methods/classes
> * Fix sanity check --no* BoolOpts
> * Add extension requirement in port-security api test
> * Fix for adding gateway with IP outside subnet
> * Add the rebinding chance in _bind_port_if_needed
> * DHCP: release DHCP port if not enough memory
> * DHCP: fix regression with DNS nameservers
> * DHCP: handle advertise_mtu=True when plugin does not set mtu values
> * Disable IPv6 on bridge devices in LinuxBridgeManager
> * ML2: delete_port on deadlock during binding
> * ML2: Add tests to validate quota usage tracking
> * Postpone heavy policy check for ports to later
> * Static routes not added to qrouter namespace for DVR
> * Make add_tap_interface resilient to removal
> * Fix bug when enable configuration named dnsmasq_base_log_dir
> * Wait for the watch process in test case
> * Trigger dhcp port_update for new auto_address subnets
> * Add generated port id to port dict
> * Protect 'show' and 'index' with Retry decorator

* Add unit test cases for linuxbridge agent when prevent_arp_spoofing is True
* Rule, member updates are missed with enhanced rpc
* Add relationship between port and floating ip
* OVS agent should fail if it cannot get DVR mac address
* DVR: Optimize check_ports_exist_on_l3_agent()
* DVR: When updating port's fixed_ips, update arp
* DVR: Fix _notify_l3_agent_new_port for proper arp update
* DVR: Notify specific agent when deleting floating ip
* DVR: Handle dvr serviceable port's host change
* DVR: Notify specific agent when creating floating ip
* DVR: Only notify needed agents on new VM port creation
* DVR: Do not reschedule the l3 agent running on compute node
* Change check_ports_exist_on_l3agent to pass the subnet_ids
* Add systemd notification after reporting initial state
* Raise RetryRequest on policy parent not found
* Keep reading stdout/stderr until after kill
* Revert "Revert "Revert "Remove TEMPEST_CONFIG_DIR in the api tox env"""
* Ensure that tunnels are fully reset on ovs restart
* Update HA router state if agent is not active
* Resync L3, DHCP and OVS/LB agents upon revival
* Fix floatingip status for an HA router
* Fix L3 HA with IPv6
* Make object creation methods in l3_hamode_db atomic
* Cache the ARP entries in L3 Agent for DVR
* Cleanup veth-pairs in default netns for functional tests
* Do not prohibit VXLAN over IPv6
* Remove 'validate' key in 'type:dict_or_nodata' type
* Fix get_subnet_for_dvr() to return correct gateway mac
* Check missed ip6tables utility
* SR-IOV: Fix macvtap assigned vf check when kernel < 3.13
* Make security_groups_provider_updated work with Kilo agents
* Imported Translations from Zanata
* Revert "Change function call order in ovs_neutron_agent."
* Remove check on DHCP enabled subnets while scheduling dvr
* Check gateway IP address when updating subnet
* Add tests that constrain database query count
* Do not call add_ha_port inside a transaction
* Log INFO message when setting admin state up flag to False for OVS port
* Call _allocate_vr_id outside of transaction
* Move notifications before database retry decorator
* Imported translations from Zanata
* Run functional gate jobs in a constrained environment
* Tox: Remove fullstack env, keep only dsvm-fullstack
* Force L3 agent to resynchronize routers that it could not configure
* Support migration of legacy routers to HA and back
* Catch known exceptions when deleting last HA router
* test_migrations: Avoid returning a filter object for python3
* move usage_audit to cmd/eventlet package
* Do not autoreschedule routers if l3 agent is back online
* Make port binding message on dead agents clear
* Disallow updating SG rule direction in RESOURCE_ATTRIBUTE_MAP
* Force service provider relationships to load
* Avoid full_sync in l3_agent for router updates

* In port_dead, handle case when port already deleted
* Kill the vrrp orphan process when (re)spawn keepalived
* Add check that list of agents is not empty in _get_enabled_agents
* Batch db segment retrieval
* Ignore possible suffix in iproute commands.
* Add compatibility with iproute2 >= 4.0
* Tune _get_candidates for faster scheduling in dvr
* Separate rbac calculation from _make_network_dict
* Skip keepalived_respawns test
* Support Unicode request_id on Python 3
* Validate local_ip for linuxbridge-agent
* Use diffs for iptables restore instead of all rules
* Fix time stamp in RBAC extension
* Notify about port create/update unconditionally
* Ensure l3 agent receives notification about added router
* get_device_by_ip: don't fail if device was deleted
* Make fullstack test_connectivity tests more forgiving
* Adding security-groups unit tests
* Check missed IPSet utility using neutron-sanity-check
* Remove duplicate deprecation messages for quota_items option
* Lower l2pop "isn't bound to any segment" log to debug

APPENDIX A. REVISION HISTORY

Revision 8.0.0-0   **Wed Feb 3 2016**   **Red Hat OpenStack Platform Docs Team**

Initial revision for Red Hat OpenStack Platform 8.0