



Red Hat OpenStack Platform 9

Release Notes

Release details for Red Hat OpenStack Platform 9

Red Hat OpenStack Platform 9 Release Notes

Release details for Red Hat OpenStack Platform 9

OpenStack Documentation Team

Red Hat Customer Content Services

rhos-docs@redhat.com

Legal Notice

Copyright © 2016 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

CHAPTER 1. INTRODUCTION	3
1.1. ABOUT THIS RELEASE	3
1.2. REQUIREMENTS	3
1.3. DEPLOYMENT LIMITS	4
1.4. DATABASE SIZE MANAGEMENT	4
1.5. CERTIFIED DRIVERS AND PLUG-INS	4
1.6. CERTIFIED GUEST OPERATING SYSTEMS	4
1.7. HYPERVISOR SUPPORT	4
1.8. CONTENT DELIVERY NETWORK (CDN) CHANNELS	5
1.9. PRODUCT SUPPORT	6
CHAPTER 2. TOP NEW FEATURES	7
2.1. RED HAT OPENSTACK PLATFORM DIRECTOR	7
2.2. BLOCK STORAGE	7
2.3. COMPUTE	8
2.4. IDENTITY	8
2.5. IMAGE SERVICE	8
2.6. OPENSTACK NETWORKING	9
2.7. TELEMETRY	9
2.8. HIGH AVAILABILITY	9
2.9. OTHER FEATURES	10
2.10. TECHNOLOGY PREVIEWS	10
CHAPTER 3. RELEASE INFORMATION	14
3.1. ENHANCEMENTS	14
3.2. KNOWN ISSUES	18
3.3. DEPRECATED FUNCTIONALITY	19
CHAPTER 4. TECHNICAL NOTES	20
4.1. RHEA-2016:1597 – RED HAT OPENSTACK PLATFORM 9 RELEASE CANDIDATE ADVISORY	20
4.2. RHEA-2016:1599 – RED HAT OPENSTACK PLATFORM 9 DIRECTOR RELEASE CANDIDATE ADVISORY	30

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 Mitaka, 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 "Mitaka" 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 "Mitaka" release itself are available at the following location:

<https://wiki.openstack.org/wiki/ReleaseNotes/Mitaka>

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 9.

You can install Red Hat OpenStack Platform 9 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 9 for RHEL 7 (RPMS)	rhel-7-server-openstack-9-rpms
Red Hat OpenStack Platform 9 director for RHEL 7 (RPMS)	rhel-7-server-openstack-9-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 9 Operational Tools for RHEL 7 (RPMS)	rhel-7-server-openstack-9-optools-rpms

Channels to Disable

The following table outlines the channels you must disable to ensure Red Hat OpenStack Platform 9 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

This section outlines the top new features for the director.

Changes to OpenStack Telemetry (ceilometer) on Overcloud Deployments

Overcloud deployed with Red Hat OpenStack Platform 9 director now use new components for OpenStack Telemetry Metrics (gnocchi) and OpenStack Telemetry Alarming (aodh).

Changes to OpenStack Identity (keystone) on Overcloud Deployments

Overcloud deployed with Red Hat OpenStack Platform 9 director now configures OpenStack Identity (keystone) as a WSGI application under ``httpd`` instead of a standalone service. This change aims to enhance the security of the service.

OpenStack Clustering (sahara) included in Overcloud Deployments

Overcloud deployed with Red Hat OpenStack Platform 9 director now includes OpenStack Clustering (sahara).

Upgrade Overcloud from Red Hat OpenStack Platform 8 to 9

The Red Hat OpenStack Platform director provides the ability to upgrade your Overcloud to Red Hat OpenStack Platform 9. This includes installing new components such as OpenStack Telemetry Metrics (gnocchi), OpenStack Telemetry Alarming (aodh), and OpenStack Clustering (sahara). The upgrade process also modifies the OpenStack Identity (keystone) service to run as a WSGI application under ``httpd`` instead of a standalone service.

Backwards Compatibility for Red Hat OpenStack 8 Overclouds

The Red Hat OpenStack Platform 9 director can manage Overclouds that use Red Hat OpenStack Platform 8. This means you can upgrade the Undercloud host to the latest version and still manage an Overcloud of the previous version. This helps support the lifecycle for the director, which is shorter than the lifecycle of the core Red Hat OpenStack Platform product.

2.2. BLOCK STORAGE

This section outlines the top new features for the Block storage service.

Snapshot Backups

You can now back up snapshots, providing another layer of data protection. This allows you to protect the snapshots taken from the volumes on a backup device, separately from the storage back end.

Improved Volume Replication (API v2.1)

The replication API now allows replication on a per-volume basis. Now, an entire storage back end can be configured to failover all of its hosted volumes to a secondary device in the event of hardware failure.

Volume Encryption Through Dashboard

Volume Encryption can now be managed through the dashboard.

2.3. COMPUTE

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

Libvirt Hardware Policy from libosinfo

Optimized guest performance using the operating system information database (libosinfo), which contains metadata about operating systems and the virtual hardware they support. Integration with libosinfo allows users to set the `os_name` image property to a valid short ID such as `rhel7` or `winxp` or URI for the operating system. The Compute service then automatically determines other optimal properties for the guest operating system, reducing the number of properties that must be manually set for each guest.

Performance-Based Thread Placement Policies for NFV and HPC Workloads

Earlier releases added support for instances with dedicated CPU resources and NUMA topology awareness that default to preferring the use of sibling threads for vCPUs where available using Simultaneous multithreading (SMT). With this release, this behavior can now be configured using image properties and flavor extra specifications:

- **prefer** - Default, prefers placing guest vCPUs on sibling threads where they are available. Host may or may not have SMT support.
- **isolate** - Isolates threads placing guest vCPUs on different physical cores. On systems with SMT support ensures that no vCPUs from other guests are placed on those cores.
- **require** - Requires the use of thread siblings. The host must have SMT support.

2.4. IDENTITY

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

Federation

This release adds support for federation, allowing you to configure your Red Hat OpenStack Platform environment to allow the credentials of existing authentication providers to access resources in that environment.

Identity Service in Apache HTTPD

With this update, the Red Hat OpenStack Platform director now configures the Identity service to run under Apache using WSGI. As a result, the Identity service now runs under the Apache HTTPD service instead of as an eventlet.

2.5. IMAGE SERVICE

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

Volume Back End Upload/Download

Images can now be uploaded to and downloaded from Block Storage volumes in the same manner as Object Storage stores.

2.6. OPENSTACK NETWORKING

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

Role-based Access Control (RBAC) for QoS policies

This release adds Role-based Access Control (RBAC) for QoS policies. As a result, you can now apply QoS policies to certain projects. For example, you can now create a QoS policy that allows for lower-priority network traffic, and have it only apply to certain projects.

RBAC for External Networks

External networks can now be controlled using the RBAC framework. This allows networks to be made available to specific tenants (as opposed to all tenants), which can then be used as an external gateway for routers and floating IPs.

Purge a Project's Networking

Previously, after deleting a project, you might have noticed the presence of stale resources that were once allocated to the project. This included networks, routers, and ports. Previously, these stale resources had to be manually deleted, while also being mindful of deleting them in the correct order. This has been addressed in Red Hat OpenStack Platform 9, where you can instead use the `neutron purge` command to delete all the neutron resources that once belonged to a particular project.

Timestamp Fields

Timestamp fields are now added to the neutron core resources: network, subnet, and ports. This enables more powerful system monitoring options, allowing you to now query neutron information against a specific period of time.

Description Fields

Optional description fields have been added to security group rules, networks, ports, routers, and floating IPs. This allows users to store descriptive information about these entities.

2.7. TELEMETRY

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

Gnocchi Integration

Gnocchi is now fully integrated with ceilometer, and is fully supported as a back end. Configuring a gnocchi back end will improve ceilometer's response time when accessing data.

Aodh Replaces ceilometer Alarm

Alarming is now using the discrete Aodh project. Deployments upgrading from Red Hat OpenStack Platform 8 to 9 will be migrated automatically.

2.8. HIGH AVAILABILITY

This section outlines the top new features for high availability.

Keystone Constraints Removed

Pacemaker constraints for keystone have been removed in this release. As a result, keystone can

now be started, stopped, or restarted without affecting the availability of other OpenStack services. Also, OpenStack services can now be started, stopped, or restarted independently of the status of the keystone service.

2.9. OTHER FEATURES

Update Container on Fast-POST (Object Storage)

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

Updated Plug-In Support and Improved Deployment (Data Processing)

The OpenStack Data Processing service (`sahara`) can now be deployed through director for quick, easy deployments and automatic upgrades. In addition, the service now supports the CDH 5.5.0 plug-in.

Improved Usability for High Availability

Pacemaker constraints on the Identity service have been removed. The Identity service can now start, stop, and restart without affecting the availability of other OpenStack services. In addition, OpenStack services can also start, stop, and restart independent of the availability of the Identity service.

2.10. TECHNOLOGY PREVIEWS

This section outlines features that are in technology preview in Red Hat OpenStack Platform 9.



NOTE

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

2.10.1. New Technology Previews

The following new features are provided as technology previews:

Google Cloud storage backup driver (Block Storage)

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

Shared File System Service

The Shared File System service (`manila`) is still included as a Technology Preview. With this release, you can test the service with the following drivers:

- NetApp (`manila.share.drivers.netapp.common.NetAppDriver`)
- CephFS native driver (`manila.share.drivers.cephfs.cephfs_native.CephFSNativeDriver`)

The new CephFS native driver allows the Shared File System service to export shared CephFS file systems to guests through the Ceph network protocol. Instances must have a Ceph client installed

to mount the file system. The CephFS file system is included in Red Hat Ceph Storage 2.0 as a technology preview as well.

At-Rest Encryption (Object Storage)

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

OpenDaylight Beryllium SR2

OpenDaylight Beryllium SR2 is now available on this release as a Technology Preview.

Red Hat SSO

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

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

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.

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.

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.

CHAPTER 3. RELEASE INFORMATION

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

Notes for updates released during the support lifecycle of this Red Hat OpenStack Platform release will appear in the advisory text associated with each update or the *Red Hat OpenStack Platform Technical Notes*. This document is available from the following page:

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

3.1. ENHANCEMENTS

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

BZ#1238592

Previously, the "nova list" command displayed instances as running when a compute node failed. Now, the instance state is updated when the hosting compute node is down. As a result, users can trust the "nova list" output for uptime monitoring.

BZ#1183796

RBD snapshots and cloning are now used for Ceph-based ephemeral disk snapshots. With this update, data is manipulated within the Ceph server, rather than transferred across nodes, resulting in better snapshotting performance for Ceph.

BZ#1316599

This enhancement adds the ability to tell the container or account server to reverse the object listings. This capability allows you to break out versioned objects in middleware. As a result, the internal architecture is reorganized for safety reasons; in addition, reverse listings are available to client applications, if needed.

BZ#1316594

This enhancement adds improved replica placement, and protection from duplicated assignments. This was added because, in the traditional Swift layout, the act of accidentally assigning two replicas of a partition to the same device resulted in a silent reduction of durability. As a result, duplicate assignments are prevented, thereby adhering to calculated guarantees. However, since this requires the number of devices to be no less than the number of replicas, it is possible for certain incorrect old rings to be considered invalid. Consequently, it is still possible to have the number of zones be smaller than the number of replicas.

BZ#1300417

With this update, a new enhancement adds a new boolean property 'router_external' to the 'OS::Neutron::ProviderNet' resource. This option allows the template author to specify whether the network contains an external router.

BZ#1170372

With this update, the eventlet system for keystone has been deprecated upstream. Red Hat OpenStack Platform director now configures keystone to run under apache using WSGI. This change was due to the Keystone project's recommendation that keystone deployment occurs within WSGI. As a result, the keystone service now runs under the apache httpd service.

BZ#1325673

This update adds Role-based Access Control (RBAC) for QoS policies. As a result, you can now apply QoS policies to certain projects. For example, you can now create a QoS policy that allows for lower-priority network traffic, and have it only apply to certain projects.

BZ#1327866

With this update, a new enhancement allows finding a suitable disk device for the Bare Metal Provisioning service to deploy the image onto by using its name or path. Some devices have persistent names (for example, RAID) and this new feature allows operators to use these names instead of having to use the disk WWN, serial, model names and so on.

BZ#1337755

This enhancement adds support for in-band cleaning of the iSCSI drivers. Cleaning steps, such as disk erase, and in-band RAID configuration, among others, can now be performed on the nodes using the drivers. The running of cleaning steps allows for improved security when recycling the nodes in ironic, allowing you to erase all the data from previous tenants and/or run checks to see if the machine wasn't compromised.

As a result, drivers, such as pxe_ipmitool, pxe_drac, pxe_iboot, pxe_ilo, pxe_amt, pxe_wol, among others, can now run in-band cleaning steps.

BZ#1339762

This update provides a client for the Aodh API. The client consists of a Python API, which is available in the "aodhclient" module, and a command-line script, which is installed as the "aodh" command. Both the Python API and the command-line script implement the entire Aodh API.

BZ#1348905

With this enhancement, the act of evacuating instances with pinned CPUs can result in these instances being hosted on a hypervisor which already handles instances with the same pinning configuration.

This was added because the resource tracker does not track CPU pinning for instances on hosts.

As a result, a condition has been added to the NUMATopologyFilter filter, which passes on hosts which already manage an instance with same CPU pinning configuration as the instance being evacuated.

BZ#1348606

The python-wsgi_intercept package has been added to Red Hat OpenStack Platform 9.

This package is an install dependency for python-gabbi, which in turn is required by openstack-gnocchi.

As a result, openstack-gnocchi and python-gabbi install without dependency errors.

BZ#1337648

CDH version 5.5 is now available in the packaged version of the CDH plugin, and is enabled by default.

BZ#1334469

This release now supports the 'map_merge' function in heat. This function allows users to merge maps together and makes values in latter maps override those in earlier ones. This can be useful when composing maps containing configuration data into a single, consolidated map.

For more information, see

http://docs.openstack.org/developer/heat/template_guide/hot_spec.html#map-merge.

BZ#1365175

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

BZ#1337762

This enhancement allows agent drivers (drivers prefixed with "agent_") to deploy partition images.

This capability was added because all drivers in OpenStack bare metal provisioning (ironic) should be able to deploy full disk images (these are images that contain a partition table with a bootloader, among other properties). This also includes partition images, which are images with

a root filesystem.

As a result, agent drivers are now able to deploy partition images, as well as full disk images.

BZ#1347347

The OpenStack data processing service (sahara) is now deployable via OSP-d. The current release will deploy Sahara to all controller nodes and make it available for use. As this is the first release to include Sahara integration with OSP director, the service is minimally configurable, but also places no additional burden on the installer; Sahara will be enabled on overcloud controller nodes by default.

BZ#1228451

With this update, an availability zone groups network nodes that run services such as DHCP and routers. It is defined as an agent's attribute on the network node, and is used to make network resources highly available. The operators group the nodes that are attached to different power sources under separate availability zones and configure scheduling for resources with high availability so that they are scheduled to different availability zones. This allows users to associate an availability zone with their routers and networks to ensure that the risk is spread across different zones.

BZ#1348609

The python-colorama package has been added to Red Hat OpenStack Platform 9.

This package is an install dependency for python-gabbi, which in turn is required by openstack-gnocchi.

As a result, openstack-gnocchi and python-gabbi install without dependency errors.

BZ#1337739

This enhancement adds manual cleaning, which allows operators to move a node directly into a cleaning state, from a manageable state.

This was added because operators may run cleaning steps for various reasons, including: Building RAID, erasing devices, among others.

As a result, operators are now able to use the OpenStack Bare Metal (ironic) API to manually start the cleaning process for the ironic nodes, choosing exactly which steps should be run.

BZ#1334467

With this enhancement, you can mark a resource as unhealthy using an API call in Orchestration (heat), so that it will be replaced on a subsequent stack update.

For example, if a server in a `ResourceGroup` fails, rather than blacklisting (which would mean non-contiguous indices and any replacement getting a new name) the server can be marked unhealthy so that Heat will replace it with a new one with the same name and index in

the group.

The new `mark-unhealthy` command allows a resource to be put in the `CHECK_FAILED` state. Any stack updates subsequently performed will replace the unhealthy resource (as they do with all resources in an `*_FAILED` state).

BZ#1334468

Using the resource registry in the environment, a user can set a hook which will pause delete actions on these resources. This allows users to take specific actions when a resource is deleted, and perform extra validation when critical elements are removed. As a result, when a resource with a pre-delete hook is about to be deleted, Heat will pause until the resource is signaled with `{'unset_hook': 'pre-delete'}` as data.

BZ#1334463

When multiple environment files are specified, they are combined in the engine instead of the client. This provides heat enough information to correctly orchestrate a stack.

3.2. KNOWN ISSUES

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

BZ#1226859

Scaling up from 0 Ceph Storage nodes is not supported

BZ#1336237

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

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

BZ#1362528

At current, `dhcpv6` on interfaces cannot be disabled by using `nic` templates, making it impossible to control IP assignment and IPv6 routes on interfaces connected to a network that runs a `dhcpv6` server. As a workaround, add the following to `/etc/sysconfig/network-scripts/ifcfg-eth4`:

```
IPV6_AUTOCONF=no
```

BZ#1463060

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.

3.3. DEPRECATED FUNCTIONALITY

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

BZ#1373985

In accordance with the upstream project, the LBaaS v1 API is being deprecated with Red Hat OpenStack Platform 9 and is expected to be removed with Red Hat OpenStack Platform 10. New customers are encouraged to use the LBaaS v2 API.

BZ#1346936

Based on customer feedback, support for a commercially available database is an absolute must for an enterprise cloud. As such, we have decided to concentrate our effort in building successful partnerships to meet this demand rather than building a fully open-source solution that might not benefit many of our customers.

In line with this, the OpenStack Trove service (which has been a Technology Preview feature so far) will no longer be included in Red Hat OpenStack Platform 10 and up. Instead, we are working with a trusted partner to provide customers with a production-ready DBaaS service. Please contact your sales account manager to learn more about this option.

BZ#1341838

Old versions of the `puppet-ceph-external.yaml` Heat environment file are not supported anymore. A newer version of this file is shipped with the 9.0 templates and must be used instead. Any customizations made to the old `puppet-ceph-external.yaml` file must be ported into the newer version.

CHAPTER 4. TECHNICAL NOTES

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

4.1. RHEA-2016:1597 – RED HAT OPENSTACK PLATFORM 9 RELEASE CANDIDATE ADVISORY

The bugs contained in this section are addressed by advisory RHEA-2016:1597. Further information about this advisory is available at <https://access.redhat.com/errata/RHEA-2016:1597.html>.

4.1.1. General

BZ#1341486

With this update, the LBaaS dashboard was moved out of horizon, and is now a separate plugin.

As a result, you can install the LBaaS dashboard using ``yum install neutron-lbaas-ui``. You will need to restart httpd to apply the change.

4.1.2. keycloak-httpd-client-install

BZ#1343228

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

4.1.3. mariadb-galera

BZ#1346067

Previously, the RPM for ``mariadb-galera`` included a step to generate TLS certificates for use in Galera SSL communication. However, when the installed RPMs were used with containers that were then replicated, the TLS certificates themselves would be replicated as well. Consequently, copies of a container would contain a TLS certificate identical to the original, creating a security condition if these certificates were actually used.

With this update, the RPM package no longer generates the certificates. As a result, no certificate is generated which may be present in a container. Certificates can be generated manually if SSL configuration of Galera is needed. Note that Red Hat OpenStack director currently does not configure Galera for SSL.

4.1.4. opendaylight

BZ#1362605

OpenDaylight Beryllium SR2 is now available on this release as a Technology Preview.

4.1.5. openstack-aodh

BZ#1336053

This rebase package includes aodh updates under version openstack-aodh-2.0.1-3.el7ost.

For the full list of changes, please refer to the upstream release notes: <http://docs.openstack.org/releasenotes/aodh/mitaka.html#id2>

BZ#1339762

This update provides a client for the Aodh API. The client consists of a Python API, which is available in the "aodhclient" module, and a command-line script, which is installed as the "aodh" command. Both the Python API and the command-line script implement the entire Aodh API.

BZ#1341764

This aodh rebase package includes a notable fix under version openstack-aodh-2.0.1-1.el7ost:

* Bug 1575530 - Update was added to fix and improve the partition coordinator, making sure that input tasks can be correctly distributed to partition members.

BZ#1353519

Previously, when creating an alarm of type ``gnocchi_aggregation_by_metrics_threshold``, the evaluator would throw an error, causing an exception.

This update addresses this issue by setting ``needed_overlap`` to always apply on aggregation.

BZ#1357880

Prior to this update, composite alarms continued to send notifications on each refresh. This resulted in unnecessary notifications appearing in the log files.

This update addresses this by no longer sending notifications on each alarm refresh.

4.1.6. openstack-ceilometer

BZ#1304982

Previously, `gnocchi-dispatcher` options were not present in

``ceilometer.conf`` by default. Consequently, the options had to be manually added. With this update, the `gnocchi-dispatcher` options are included by default when `ceilometer` is installed.

BZ#1337961

Prior to this update, events could not be filtered by time-based fields. Consequently, ``le`` and ``ge`` queries did not work on time-based fields. This update adds new query operators, with the result that the ``le`` and ``ge`` operators should now work on time-based queries.

BZ#1349350

Telemetry (`ceilometer`) `dbsync` creates an old alarming table, as a result of still running the ``sqlalchemy-migrate`` code. Consequently, `Aodh` `dbsync` sees no reference to `Alembic`, and has `SQLAlchemy` create the necessary tables. However, the tables are already present, so nothing is done, and the database is stamped to the latest version of `Aodh` `Alembic`. You can avoid this issue by not creating alarm tables in `ceilometer` `dbsync`.

4.1.7. openstack-cinder

BZ#1312944

This update adds RPM requirements to ensure that the required python libraries for the Google Cloud backup driver are present.

BZ#1333547

Previously, `python-taskflow` was missing a dependency on a suitable version of `python-networkx`. Consequently, ``cinder create volume`` did not function as expected. With this update, `python-taskflow` package has the correct dependencies, and ``cinder create volume`` works as expected.

4.1.8. openstack-gnocchi

BZ#1341704

This `openstack-gnocchi` rebase package adds notable updates under version: `openstack-gnocchi-2.1.1-1.el7ost`
For more information, see the upstream release milestone:
<https://launchpad.net/gnocchi/+milestone/2.1.1>

4.1.9. openstack-heat

BZ#1334463

■

When multiple environment files are specified, they are combined in the engine instead of the client. This provides heat enough information to correctly orchestrate a stack.

BZ#1334468

Using the resource registry in the environment, a user can set a hook which will pause delete actions on these resources. This allows users to take specific actions when a resource is deleted, and perform extra validation when critical elements are removed. As a result, when a resource with a pre-delete hook is about to be deleted, Heat will pause until the resource is signaled with `{'unset_hook': 'pre-delete'}` as data.

4.1.10. openstack-ironic**BZ#1291382**

Previously, the iPXE driver had a conditional that prevented nodes from being configured with UEFI boot mode. Consequently, iPXE driver users could not configure their nodes with UEFI and were forced to use the BIOS instead.

With this update, the conditional was removed, and users of the iPXE driver can now deploy their nodes in UEFI mode.

BZ#1337739

This enhancement adds manual cleaning, which allows operators to move a node directly into a cleaning state, from a manageable state.

This was added because operators may run cleaning steps for various reasons, including: Building RAID, erasing devices, among others.

As a result, operators are now able to use the OpenStack Bare Metal (ironic) API to manually start the cleaning process for the ironic nodes, choosing exactly which steps should be run.

BZ#1337755

This enhancement adds support for in-band cleaning of the iSCSI drivers. Cleaning steps, such as disk erase, and in-band RAID configuration,

among others, can now be performed on the nodes using the drivers.

The running of cleaning steps allows for improved security when recycling the nodes in ironic, allowing you to erase all the data from previous tenants and/or run checks to see if the machine wasn't compromised.

As a result, drivers, such as `pxe_ipmitool`, `pxe_drac`, `pxe_iboot`, `pxe_ilo`, `pxe_amt`, `pxe_wol`, among others, can now run in-band cleaning steps.

4.1.11. openstack-neutron**BZ#1328773**

Previously, `ipset` was not declared as a dependency of the Open vSwitch (OVS) and Linux Bridge neutron agents. However, `ipset` is a dependency of the `openstack-neutron` package, and this would result in nodes that had the packages for the Open vSwitch or Linux Bridge agent installed, but did not have `openstack-neutron` installed. The L2 agents required `ipset` for configuration of security groups.

With this update, `ipset` is a dependency of the `openstack-openvswitch-agent`, and the `openstack-linuxbridge-agent` packages depend on `ipset`.

BZ#1328781

Previously, the `openstack-neutron-common` package did not require the `shadow-utils` package. This prevented the 'neutron' user from being created when the `openstack-neutron-common` package was installed, resulting in neutron being unable to execute commands on hypervisors. With this update, the `openstack-neutron-common` package now requires the `shadow-utils` package, and the 'neutron' user is correctly created.

BZ#1346822

Previously, when bridge ports were missing for `br-int` and `br-tun` during agent startup, it checked for the patch ports `int-br-ex` and `phy-br-ex` before adding them. However, the function used to check their existence was `get_port_ofport()`, which retried the check because of the `@_ofport_retry` decoration.

Consequently, this caused the restart to become unnecessarily slow because of the retries.

With this update, the existence of ports is checked with `port_exists()` instead of `get_port_ofport()`. As a result, no slowdown occurs on startup when the bridge ports are missing for `br-int` and `br-tun`.

4.1.12. openstack-nova

BZ#1328957

With this update, the `openstack-nova` package is now re-based to upstream version 13.1.0.

BZ#1331420

Previously, when booting instances, the nova API automatically added a default security group if nothing was specified, which should not be done on a network with option `'port_security_enabled=False'`

Consequently, the boot process would fail for users booting an instance that was attached to a network with port security disabled.

With this update, nova no longer adds a default security group to a port created for an instance on a network with `port_security_enabled=False`

As a result, the boot process works as expected, and the port attached to the instance does not have a default security group attached.

NOTE: a known bug in the dashboard still indicates that a default security group is attached to the instance, but this only occurs during the first attempt at booting the instance.

BZ#1332599

With `serial_console` enabled, repeatedly starting and stopping an instance eventually caused the compute service to run out of ports. When this occurred, attempting to start an instance resulted in a `'SocketPortRangeExhaustedException'` error, preventing instances from being started. This was because while the compute service creates a port when an instance is started (as opposed to created), it only releases a port when an instance is deleted (as opposed to stopped).

In this update, the method that destroys the guest from a libvirt perspective now also releases serial ports. This ensures that a serial port is released as soon as it is no longer required by an instance, thereby making the port available again.

BZ#1335835

When creating snapshots, the compute API now omits disk format and container format details from the image API request. This helps ensure that the driver will use the correct snapshot image format. This, in turn, prevents snapshots from failing with a `BadRequest` if its image is converted to a format other than what the base image uses.

BZ#1341612

Previously, under BZ#1332599, a regression was discovered which caused an instance to lose its serial ports after a hard-reboot. Consequently, it was not possible to connect to the instance through a serial console. This issue occurred during a hard-reboot process, because the serial ports on a host were released, but since the domain XML was still defining them, the process to acquire ports on the host during the boot was not executed.

To address this issue, nova will now undefine the domain from libvirt after having destroyed it.

BZ#1342578

When assigning an SR-IOV Virtual Function (VF) device to an instance, its corresponding Physical Function (PF) device is correctly masked as unavailable in the database. However, in past releases, deleting the instance did not update the PF as available. As a result, PCI devices were never released from the database after instances which used them were deleted.

With this update, nova now maintains an in-memory tree of PCI devices, then periodically flushes it into the database. This helps update the database on information about available devices.

BZ#1342953

In previous releases, the systemd unit file for openstack-nova-compute was missing a required dependency on libvirtd. As such, restarting openstack-nova-compute failed whenever libvirtd was not already running. This update adds the dependency.

BZ#1348905

With this enhancement, the act of evacuating instances with pinned CPUs can result in these instances being hosted on a hypervisor which already handles instances with the same pinning configuration.

This was added because the resource tracker does not track CPU pinning for instances on hosts.

As a result, a condition has been added to the NUMATopologyFilter filter, which passes on hosts which already manage an instance with same CPU pinning configuration as the instance being evacuated.

4.1.13. openstack-packstack

BZ#1332525

Previously, Packstack tried to create a gnocchi database when ceilometer installation was disabled. As a result, disabling ceilometer caused Packstack installations to fail, as some parameters required for creating a gnocchi database were not passed. With this release, Packstack no longer attempts to create a gnocchi database if ceilometer is disabled.

4.1.14. openstack-puppet-modules

BZ#1337250

Previously, iPXE would freeze during HTTP download resulting in the Bare Metal Provisioning (ironic) service to hang.

This update makes sure that iPXE retries to boot from the network in case of a failure. The '--timeout' option can be used to avoid an unlimited freeze. As a result, a freeze does not occur during the HTTP download.

BZ#1349891

The service providers responsible for configuring VPNaaS and LBaaS created additional files that were never included in the service startup. This prevented neutron-server from starting if LBaaS was enabled.

Previously, the LBaaS service config provider was updated to put service providers directly into /etc/neutron/neutron.conf. Updating VPNaaS to do the same would have caused it to overwrite the 'service_provider' value set by LBaaS, or vice-versa. So to address this, this update moves the 'neutron_config' provider from ini_setting to openstackconfig and adds a

variable to `neutron::server` to manage service providers. This, in turn, prevents VPNaaS and LBaaS from overwriting each others' `service_provider` values. As a result, enabling LBaaS no longer prevents `neutron-server` from starting.

BZ#1353971

The Ceph puppet module (`puppet-ceph`) did not update CephX keyrings when the `caps` parameter for a key were changed. This caused overcloud upgrades to fail, as caps to operate on the new 'metrics' pool were not added to the CephX keyring.

With this update, `puppet-ceph` regenerates the `virsh` secret or updates its key if either `'rbd_keyring'` or `'linvirt_rbd_secret_uuid'` change. This ensures that the CephX keyring is updated as required when secrets or caps change.

4.1.15. openstack-selinux

BZ#1315457

Previously, the absence of SELinux policy that allowed the Compute API to be started in WSGI with Apache resulted in an AVC in the `audit.log`.

With this update, Compute is able to bond to the HTTP's port and runs without errors when started in WSGI with Apache.

BZ#1325623

Previously, running the Block Storage API in WSGI with Apache and SELinux in the 'enforce' mode resulted in an AVC, as SELinux prevented the `'/usr/sbin/httpd'` from access to the `'/var/log/cinder/cinder-api.log'` file.

With this update, `'httpd'` is allowed access to the Block Storage API log file. As a result, the Block Storage API in WSGI runs without AVCs.

4.1.16. openstack-swift

BZ#1316594

This enhancement adds improved replica placement, and protection from duplicated assignments.

This was added because, in the traditional Swift layout, the act of accidentally assigning two replicas of a partition to the same device resulted in a silent reduction of durability.

As a result, duplicate assignments are prevented, thereby adhering to calculated guarantees. However, since this requires the number of devices to be no less than the number of replicas, it is possible for

certain incorrect old rings to be considered invalid. Consequently, it is still possible to have the number of zones be smaller than the number of replicas.

BZ#1316599

This enhancement adds the ability to tell the container or account server to reverse the object listings. This capability allows you to break out versioned objects in middleware.

As a result, the internal architecture is reorganized for safety reasons; in addition, reverse listings are available to client applications, if needed.

4.1.17. python-ceilometerclient

BZ#1315324

Alarms created with the type 'gnocchi-resource-threshold' had empty fields for the 'project_id' and 'user_id' fields.

With this update, when alarms are created with type 'gnocchi-resource-threshold', the 'project_id' and 'user_id' fields are populated.

4.1.18. python-colorama

BZ#1348609

The python-colorama package has been added to Red Hat OpenStack Platform 9.

This package is an install dependency for python-gabbi, which in turn is required by openstack-gnocchi.

As a result, openstack-gnocchi and python-gabbi install without dependency errors.

4.1.19. python-cradox

BZ#1343144

The Time Series Database-as-a-Service (gnocchi) Ceph storage driver had a dependency on the 'python-cradox' library, resulting in a gnocchi with a Ceph back end to fail.

With this update, the 'python-cradox' is installed by the 'openstack-puppet-modules' package. The 'python-cradox' is a Python library for the Ceph librados library which uses 'cython' instead of 'ctypes'. As a result, the Time Series Database-as-a-Service with a Ceph back end will run without errors.

4.1.20. python-django-horizon

BZ#1041987

With this update, the API part for using the Ceilometer alarm API has been added to Horizon. It enables future developments to use the API. However, at present, there is no GUI front end for this API.

BZ#1093899

This update adds support for Cinder volume encryption to Horizon, which enables administrators to manage encrypted volume types from the GUI. As a result, volume types can now be added, modified, and deleted using Horizon.

BZ#1351736

With this update, the 'python-django-horizon' packages have been rebased to version 9.0.1.

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

- * Fix workflow bug in "Create Network"
- * Fix existing metadata display in metadata widget
- * Various localisation fixes

4.1.21. python-heatclient

BZ#1281557

The 'stack-delete' command displayed information about the heat stack immediately after requesting the delete. Deleting a heat stack is an asynchronous operation so the display status may not have yet changed to 'DELETE_IN_PROGRESS', and users may find this confusing.

The 'stack-delete' command now returns nothing, which is consistent with other OpenStack services delete commands.

4.1.22. python-oslo-concurrency

BZ#1346385

This release includes updates from oslo.concurrency 3.7.1, which places new process limits required to address security vulnerability CVE-2015-5162. See <https://bugs.launchpad.net/ossa/+bug/1449062> for details about the vulnerability and fix.

4.1.23. python-wsgi_intercept

BZ#1348606

-

The `python-wsgi_intercept` package has been added to Red Hat OpenStack Platform 9.

This package is an install dependency for `python-gabbi`, which in turn is required by `openstack-gnocchi`.

As a result, `openstack-gnocchi` and `python-gabbi` install without dependency errors.

4.2. RHEA-2016:1599 – RED HAT OPENSTACK PLATFORM 9 DIRECTOR RELEASE CANDIDATE ADVISORY

The bugs contained in this section are addressed by advisory RHEA-2016:1599. Further information about this advisory is available at <https://access.redhat.com/errata/RHEA-2016:1599.html>.

4.2.1. instack-undercloud

BZ#1320816

The `ironic-inspector` was not included in the Keystone endpoint catalog in the Undercloud. This meant you could not discover the API endpoints for `ironic-inspector` through the catalog. This fix adds the `ironic-inspector` API endpoints to the Keystone catalog. You can now discover the `ironic-inspector` API endpoints through the catalog.

BZ#1341350

Slow environments experienced timeouts when Glance tried to communicate with Swift as a backend. This caused some Glance operations, such as image uploads, to fail. This fix increases the Swift proxy server's default `node_timeout` value to 60 seconds. This increases the reliability of Glance image uploads on slow environments using Swift as an image storage backend.

4.2.2. ipxe

BZ#1353361

The `ipxe-bootimages` RPM was not included in the repository for the director. This caused failure during director installation. This update adds the package to the director's repository. The RPM is now included as a part of the director's installation.

4.2.3. openstack-tripleo-heat-templates

BZ#1290121

Almost all the Overcloud Pacemaker resources depended on the Keystone resource. This meant restarting the Keystone resource after a configuration change would restart all dependent resources, which was disruptive. This fix introduces a fake `openstack-core` that the Overcloud

Pacemaker resources (including keystone) use as a dependency. This means restarting the Keystone resource no longer causes any disruption to other services.

BZ#1337511

The `ManagementNetValueSpecs` parameter used the wrong type (string) in the director's Heat templates. This caused Overcloud deployments containing the Management network to fail with the following error:

```
"Property error:
resources.ManagementNetwork.properties.ManagementNetValueSpecs: Value
must be a string".
```

This fix changes the data type of `ManagementNetValueSpecs` from string to json. The error no longer appears.

BZ#1340453

In previous releases of Heat, domain resources were created before `/etc/heat/heat.conf` was configured on the overcloud. However, domain resources depended on settings from that file; as such, these resources were not created correctly, preventing users from creating heat stacks. Users had to manually restart the Pacemaker Heat Engine resource to work around the problem.

This release corrects the sequence of steps for deploying the Heat service, thereby fixing the problem.

BZ#1341838

Old versions of the `puppet-ceph-external.yaml` Heat environment file are not supported anymore. A newer version of this file is shipped with the 9.0 templates and must be used instead. Any customizations made to the old `puppet-ceph-external.yaml` file must be ported into the newer version.

BZ#1344307

This patch updates the Help URL on the dashboard. The URL now points to the Official Red Hat OpenStack Platform documentation instead of the OpenStack upstream documentation.

BZ#1349180

The default `network-isolation.yaml` file contained values that conflicted with default `network-management.yaml` file in the director's Heat templates. If you included the `network-isolation.yaml` file after the `network-management.yaml` when creating an Overcloud, the deployment deactivates the Management network. This update refactors these files to

avoid conflicts. Now the deployment processes both the network-management.yaml and network-isolation1.yaml in any order without conflict.

BZ#1353637

The Overcloud uses a fake openstack-core resource to replace keystone (now accessed through httpd). However, the Overcloud did not create openstack-core when using an external load balancer. This caused puppet to fail when creating Pacemaker constraints that required openstack-core. This fix ensures creation of the openstack-core resource for any deployment configuration. The deployment succeeds and creates the needed constraints.

BZ#1356107

OpenStack Platform 9 deployments require an additional CephX key for the "client.openstack" user. However, the director's command line client does not generate this key for existing deployments and updates the "ceph.openstack" keyring with an empty secret. Before upgrading, generate a new CephX key and pass it to the deployment using the CephClientKey parameter in an environment file. For example:

```
parameter_defaults:
  CephClientKey: 'my_cephx_key'
```

Generate the new key with the following command:

```
$ ceph-authtool --gen-print-key
```

4.2.4. os-collect-config

BZ#1350489

The "os-collect-config" service on the Overcloud restarted on an RPM update. This caused Overcloud updates to fail. This fix changes the behavior so that "os-collect-config" does not restart on an RPM update. The Overcloud updates now succeed after an update of "os-collect-config". Note that "os-collect-config" gracefully restarts itself when "os-refresh-config" runs, so the restart on update is not required.