



Red Hat Enterprise Linux OpenStack Platform 7

Release Notes

Release details for Red Hat Enterprise Linux OpenStack Platform

Red Hat Enterprise Linux OpenStack Platform 7 Release Notes

Release details for Red Hat Enterprise Linux OpenStack Platform

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

Legal Notice

Copyright © 2015 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 Enterprise Linux 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	4
1.9. PRODUCT SUPPORT	6
1.10. PRODUCT DOCUMENTATION	7
CHAPTER 2. TOP NEW FEATURES	9
2.1. RHEL OPENSTACK PLATFORM DIRECTOR	9
2.2. BLOCK STORAGE	9
2.3. COMPUTE	10
2.4. IDENTITY	11
2.5. IMAGE SERVICE	11
2.6. OBJECT STORAGE	12
2.7. OPENSTACK INTEGRATION TEST SUITE	12
2.8. OPENSTACK NETWORKING	13
2.9. TECHNOLOGY PREVIEWS	13
CHAPTER 3. RELEASE INFORMATION	16
3.1. ENHANCEMENTS	16
3.2. TECHNOLOGY PREVIEW	33
3.3. RELEASE NOTES	33
3.4. KNOWN ISSUES	35
CHAPTER 4. UPGRADING	41
CHAPTER 5. ADDITIONAL OVERCLOUD CONFIGURATION PARAMETERS	42
CHAPTER 6. TECHNICAL NOTES	44
6.1. RHEA-2015:1548 — RED HAT ENTERPRISE LINUX OPENSTACK PLATFORM ENHANCEMENT ADVISORY	44
6.2. RHEA-2015:1549 — RED HAT ENTERPRISE LINUX OPENSTACK PLATFORM DIRECTOR RELEASE	66
6.3. RHSA-2015:1862 — MODERATE: RED HAT ENTERPRISE LINUX OPENSTACK PLATFORM 7 DIRECTOR UPDATE	87
6.4. RHSA-2015:2650 — MODERATE: RED HAT ENTERPRISE LINUX OPENSTACK PLATFORM 7 DIRECTOR UPDATE	95
6.5. RHBA-2015:2651 — RED HAT ENTERPRISE LINUX OSP 7 DIRECTOR BUG FIX ADVISORY	100
6.6. RHBA-2015:2652 — OPENSTACK-NEUTRON BUG FIX ADVISORY	104
6.7. RHBA-2016:0264 — RED HAT ENTERPRISE LINUX OSP 7 DIRECTOR BUG FIX ADVISORY	107
APPENDIX A. REVISION HISTORY	117

CHAPTER 1. INTRODUCTION

Red Hat Enterprise Linux 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 Kilo, 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 Enterprise Linux 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 Enterprise Linux OpenStack Platform is based on the OpenStack "Kilo" release. It includes additional features, known issues, and resolved issues specific to Red Hat Enterprise Linux OpenStack Platform.

Only changes specific to Red Hat Enterprise Linux OpenStack Platform are included in this document. The release notes for the OpenStack "Kilo" release itself are available at the following location:
<https://wiki.openstack.org/wiki/ReleaseNotes/Kilo>

Red Hat Enterprise Linux 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 Enterprise Linux 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 Enterprise Linux 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 Enterprise Linux OpenStack Platform: <https://access.redhat.com/site/solutions/509783>

1.2. REQUIREMENTS

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

The Red Hat Enterprise Linux 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)

1.3. DEPLOYMENT LIMITS

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

1.4. DATABASE SIZE MANAGEMENT

For recommended practices on maintaining the size of the MariaDB databases in your RHEL 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 RHEL OpenStack Platform, see [Component, Plug-In, and Driver Support in RHEL OpenStack Platform](#).

1.6. CERTIFIED GUEST OPERATING SYSTEMS

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

1.7. HYPERVISOR SUPPORT

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

With this release of Red Hat Enterprise Linux OpenStack Platform, Ironic is now fully supported. 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 Enterprise Linux OpenStack Platform 7.

You can install Red Hat Enterprise Linux OpenStack Platform 7 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 Enterprise Linux OpenStack Platform 7.0 (RPMS)	rhel-7-server-openstack-7.0-rpms
Red Hat Enterprise Linux OpenStack Platform Director 7.0 (RPMS)	rhel-7-server-openstack-7.0-director-rpms

Table 1.2. Optional Channels

Channel	Repository Name
Red Hat Enterprise Linux 7 Server - Optional	rhel-7-server-optional-rpms
Red Hat Enterprise Linux OpenStack Platform 7.0 Files	rhel-7-server-openstack-7.0-files
Red Hat Enterprise Linux OpenStack Platform 7.0 Operational Tools	rhel7-server-openstack-7.0-optools-rpms

Channels to Disable

The following table outlines the channels you must disable to ensure Red Hat Enterprise Linux OpenStack Platform 7 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 Enterprise Linux OpenStack Platform software repositories conflict with packages provided by the Extra Packages for Enterprise Linux (EPEL) software repositories. The use of Red Hat Enterprise Linux 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 Enterprise Linux OpenStack Platform.

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

- The **rhos-list** mailing list provides a forum for discussions about installing, running, and using OpenStack on Red Hat based distributions.

Subscribe at <https://www.redhat.com/mailman/listinfo/rhos-list>.

1.10. PRODUCT DOCUMENTATION

In this release of RHEL OpenStack Platform, the content published as KBase articles in previous release has been compiled into a series of guide based around the major components and administrative tasks in RHEL OpenStack platform. With this change, all documentation is now available in the **html**, **html-single**, **epub**, and **pdf** formats. Moreover, information on evaluating RHEL using the Packstack deployment tool is no longer published under the product documentation page, but several scenarios for using this tool can be found on the RHEL OpenStack Platform product page.

The following titles are included in the product documentation:

Architecture Guide

An introduction to each of the major components in RHEL OpenStack Platform, and to a set of example scenarios. This guide combines content from the *Component Overview* included in RHEL OpenStack Platform 6 with additional content that outlines sample considerations and scenarios for designing RHEL OpenStack Platform environments.

Back Up and Restore Red Hat Enterprise Linux OpenStack Platform

An introduction to backing up, restoring, and planning recovery for a RHEL OpenStack Platform director environment.

Bare Metal Provisioning

An introduction to procedures for installing, configuring, and using the Bare Metal Provisioning service in the Overcloud of a RHEL OpenStack Platform environment.

Command-Line Interface Reference

A reference to the command-line clients in RHEL OpenStack Platform, including a list of available commands, syntax, and options.

Configuration Reference

A reference to the configuration options available to each of the components in a RHEL OpenStack Platform environment and the available values.

Dell EqualLogic Back End Guide

Describes how to configure RHEL OpenStack Platform to use one or more Dell EqualLogic back ends.

Director Installation and Usage

A comprehensive guide to deploying and managing a RHEL OpenStack Platform environment using the new RHEL OpenStack Platform director.

Installation Reference

A reference that outlines the basic steps involved in manually configuring components. This guide represents a streamlined of the *Deploying OpenStack: Learning Environments (Manual Setup)* that was available in previous releases.

Introduction to the OpenStack Dashboard

An introduction to major user interface elements in the RHEL OpenStack Platform dashboard.

Instances and Images Guide

A guide to creating and managing images, and working with instances and related concepts such as volumes and containers.

Logging, Monitoring, and Troubleshooting Guide

A collection of information on the logs available in RHEL OpenStack Platform, how to configure and work with the telemetry service, methods for addressing common problems, and the Red Hat Access tab available from the dashboard.

Migrating Instances

A guide to migrating running instances and static instances between hypervisors in a RHEL OpenStack Platform environment.

Networking Guide

A comprehensive guide to networking in RHEL OpenStack Platform, including an introduction to networking, how traditional networking concepts correlate to software-defined networking, and a cookbook-style collection of advanced networking concepts, procedures, and considerations.

OpenStack Data Processing

An introduction to configuring and using the data processing service to easily provision and scale Hadoop clusters to process large datasets.

Package Manifest

A comprehensive list of the packages provided in this release of RHEL OpenStack Platform and subsequent maintenance releases.

Upgrading OpenStack

A guide to upgrading RHEL OpenStack Platform components from version 6 to version 7.

Users and Identity Management Guide

A guide to creating and working with users, projects, and roles, and integration with external directory services.

VMware Integration Guide

A outline of integration between RHEL OpenStack Platform and VMware, including how to migrate virtual machines from VMware to RHEL OpenStack Platform, integration with VMware vCenter, and integration between VMware NSX and OpenStack Networking.

CHAPTER 2. TOP NEW FEATURES

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

2.1. RHEL OPENSTACK PLATFORM DIRECTOR

The Red Hat Enterprise Linux OpenStack Platform director is a new deployment and lifecycle management tool for your OpenStack environment. It is based on the TripleO project and provides the following features:

Deployment Consistency

The Red Hat Enterprise Linux OpenStack Platform director provides a means to install an OpenStack environment using OpenStack services and API. The method involves using an underlying OpenStack instance to install and then to manage another, usually more complex, OpenStack instance through a set of image building configuration files. The director achieves this using either CLI and GUI interaction. This also involves using OpenStack's Ironic service to provision bare metal machines.

In addition to installation, configuration, and management, the director provides a set of automated benchmark and health check during and after installation, and also offers ready-state configuration for RAID, BIOS, and network interfaces.

Lifecycle Management

The Red Hat Enterprise Linux OpenStack Platform director provides tools to scale the capacity of your environment. The director can apply updates to the current version of the Red Hat Enterprise Linux OpenStack Platform environment as well as upgrade to new versions of Red Hat Enterprise Linux OpenStack Platform.

The director also integrates with other Red Hat products such as:

- Red Hat Ceph Storage
- Red Hat Satellite 6
- Other Red Hat Cloud Infrastructure products (Red Hat CloudForms and Red Hat Enterprise Virtualization)

Accelerated Releases

Red Hat aims for new Red Hat Enterprise Linux OpenStack Platform director releases every two months.

Operational Visibility (Technology Preview)

The Red Hat Enterprise Linux OpenStack Platform director acts as a central logging tool for your OpenStack environment. This also includes notification alarms for availability and performance monitoring.

2.2. BLOCK STORAGE

Incremental backup

The Block Storage service now supports incremental backups. With this feature, you can save changes to a volume since its last backup, as opposed to performing a full backup each time. This

allows your backup operations to scale better over time as the number and sizes of your volumes grow over time.

To make incremental backups possible, the Block Storage feature now also supports snapshot-based backups. Snapshot-based backups allow you to back up volumes while they are still attached to an instance.

NFS and POSIX Backups

Cinder backup now has support for using both NFS-supplied and POSIX-supplied data repositories as backup targets.

Private volume types

When creating a volume type, you can now set it to private. Private volume types are only accessible in projects where they are explicitly added. If a private volume type has no attached project, only users with the admin role can use them. Private volume types are useful for restricting access to certain volume specifications.

Enhanced iSCSI multipath support

The Block Storage service can now provide the Compute service with all available portal, IQN, and LUN information for accessing a volume. This allows the Compute service to establish multiple sessions for each target, which in turn facilitates failover in case the primary path is down.

Consistency groups

The Block Storage service now allows you to set Consistency Groups. With this, you can group multiple volumes together as a single entity; this, in turn, allows you to perform operations on multiple volumes at once, rather than individually. At this time, the only supported operation on consistency groups is snapshot creation.

Note that this capability is only available through the IBM Storwize driver.

2.3. COMPUTE

Support for quiescing file-systems during image snapshot using QEMU Guest Agent

Previously, file systems had to be quiesced manually (`fsfreeze`) before taking a snapshot of active instances for consistent backups. With this update, Compute's libvirt driver now automatically requests the QEMU Guest Agent to freeze the file systems (and applications if `fsfreeze-hook` is installed) during an image snapshot. Support for quiescing file systems enables scheduled, automatic snapshots at the block device level.

This feature is only valid if the QEMU Guest Agent is installed (`qemu-ga`) and the image metadata enables the agent (`hw_qemu_guest_agent=yes`).

OpenStack Bare Metal Provisioning (ironic)

The Bare Metal Provisioning service is now supported in Compute (using Compute and OpenStack Networking service configuration). The service integrates with the Compute service (in the same way that virtual machines are provisioned), and provides a solution for the 'bare-metal-to-trusted-tenant' use case. For example, within the OpenStack cloud:

- Hadoop clusters can be deployed on bare metal.
- Hyperscale and high-performance computing (HPC) clusters can be deployed.

- Database hosting for applications sensitive to virtual machines can be used.

The service consists of the Bare Metal Provisioning API, a Conductor, database- and hardware-specific drivers, and leverages common technologies like PXE, IPMI, and DHCP. With this release, the Bare Metal Provisioning service can also now be passed capabilities defined in a flavor's extra-spec key.

2.4. IDENTITY

Hierarchical multitenancy

Red Hat Enterprise Linux OpenStack Platform now adds support for hierarchical ownership of objects. This allows you to modify the organizational structure of RHEL OpenStack Platform, creating nested projects in Identity.

Federation with SAML

Federated Identity establishes trust between Identity Providers (IdP) and the services provided by an OpenStack Cloud to an end user. Federated Identity provides a way to securely use existing credentials to access cloud resources such as servers, volumes, and databases across multiple endpoints provided in multiple authorized clouds using a single set of credentials, without having to provision additional identities or log in multiple times. The credentials for users and groups are maintained by the user's Identity Provider.

Federated users are not mirrored in the Identity service back end (for example, using the SQL driver). The external Identity Provider is responsible for authenticating users, and communicates the result of the authentication to Identity service using SAML assertions. SAML assertion contains information about a user as provided by an Identity Provider. Identity service maps the SAML assertions to Keystone user groups and assignments created in Identity service.

Web SSO with Keystone and SAML

RHEL OpenStack Platform now provides the ability for users to authenticate via a web browser with an existing Identity Provider (IdP), through a single sign-on page.

2.5. IMAGE SERVICE

Image Conversion

Allows you to convert an image on the fly depending on the store while importing an image. (qcow/raw).

A plug-in of the import workflow provides the conversion. Based on the deployer configuration, you can either activate or deactivate this plug-in. As a deployer, you need to specify the preferred format of images for the deployment. Currently, the formats supported by qemu-img convert are **raw** and **qcow2**.

Introspection of Images

Metadata extraction allows administrators to better understand how to override certain metadata.

Several image formats include metadata inside the image data. This new feature exposes the metadata through introspection of the image. For example, you can read the metadata from a **vmdk**-formatted image to know that the disk type of the image is **"streamOptimized"**. Allowing the

Image service to perform this introspection reduces the burden placed on administrators. Exposing this metadata also helps the consumer of the image; the Compute workflow is very different based on the disk type of the image.

Currently, the metadata fields relevant to the Image service that are available are the **disk_format** and the **virtual_size**.

2.6. OBJECT STORAGE

Composite Tokens and Service Accounts

Previously, data was stored in either a dedicated service account (single project) or in the end-user's account (multi project). If data was stored in the service account, there could be issues with container or service-user deletion, and passwords and tokens were fragile. If data was stored in the end-user account, issues arose if the user picked the same name as the service, and the Object Storage service had to deal with users violating the integrity of tokens.

To solve these issues:

- Object Storage now stores service data in a separate account to the end-user's 'normal' account, but which is still linked to the end-user's project for accounting purposes. Access to the service account is managed by composite tokens.
- Composite tokens have been introduced. Composite tokens use two authentication tokens when storing data: one from the Object Storage service and one from the end user. Once the two tokens are set, the data cannot be changed without consent from both service and user. This protects data from being deleted by the end user without some kind of administrative process, and prohibits the service from making decisions on behalf of the user.

Efficient Replication for Globally Distributed Clusters

Previously, in globally distributed clusters, replicas were pushed out to each node in each region. With this update, the Object Storage service now only pushes replicas to just one remote node in each region (based on affinity); the remote node then spreads the replica out to primary nodes in the same region. By reducing the amount of data transfer between regions, this update stabilizes performance and lowers both the chance of replication delays as well as the cost of data transfers.

2.7. OPENSTACK INTEGRATION TEST SUITE

The Integration Test Suite (tempest) is now included in Red Hat Enterprise Linux OpenStack Platform. The suite offers a set of integration tests that can be run against your OpenStack environment, ensuring that your cloud works as expected. Integration testing combines and tests the individual OpenStack modules as a group, enabling you to test complete cloud functionality.

The Integration Test Suite offers the following features:

Complete testing

The suite includes API, scenario, and stress tests. The suite also includes a set of unit tests that can test the suite's code itself. You can run the entire test suite, all tests in one directory, a test module, or just a single test.

Configuration

You can manually configure the suite or use a script that pulls information from the testing environment (by querying the cloud), and which can create necessary resources or credentials.

Scalable

The suite can be run against any OpenStack cloud, regardless of the cloud's node size. It can spin up instances or volumes on each compute or storage node in the cloud, conduct tests, and then terminate them again.

Public Interfaces

The suite only runs against public interfaces (your OpenStack endpoints). No private or implementation-specific interfaces are used. Tests are not run directly against the database or hypervisors.

Choice of Authentication

Tests can be run as a regular user, a global admin user, or using some other user set of credentials.

2.8. OPENSTACK NETWORKING**Port security with ML2 and Open vSwitch**

OpenStack Networking applies anti-spoofing firewall rules by default, with the result that a VM cannot communicate using a MAC or IP address that is not configured on its network port. In Red Hat Enterprise Linux OpenStack Platform 7, it is now possible to enable or disable the security-group feature on a per port basis, using the new 'port-security-enabled' attribute. Consequently, Project administrators get granular control over the firewall's position in the network topology.

Enhancements to L3 High Availability

View the state of HA routers - Administrators are now able to view the state of High Availability routers on each node, and specifically, where the active instance is hosted. This new functionality also serves as a sanity test and offers assurance that a router is indeed active on only one node.

Support for multiple subnets on external networks - HA routers are now able to allocate floating IP addresses to all on-link subnets.

LBaaS v2 API

LBaaS version 2.0 allows for increased robustness in load-balancing deployments, including support for SSL/TLS termination. This update to v2 includes a redesign of the LBaaS architecture and the HAProxy reference plugin.

Tech Preview - DVR integration between VLANs and VXLAN/GRE

Red Hat Enterprise Linux OpenStack Platform 7.0 (kilo) adds support for interconnecting between VLAN and VXLAN/GRE when using distributed routers. This integration allows connectivity between VLANs and VXLAN/GRE tunnels in DVR.

IPv6 Support

In RHEL OpenStack Platform 7, the core OpenStack services are able to operate over IPv6 networks. At present, RHEL OpenStack Platform director will not deploy or manage nodes over IPv6-based networks.

2.9. TECHNOLOGY PREVIEWS

Distributed virtual routing, DNS-as-a-Service, and erasure coding are included as technology previews in Red Hat Enterprise Linux OpenStack Platform 7.

**NOTE**

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

Cells

OpenStack Compute includes the concept of Cells, provided by the nova-cells package, for dividing computing resources. For more information about Cells, see the [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

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 Enterprise Linux OpenStack Platform 7 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)

As a technology preview, the Object Storage service now 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 OpenStack File Share Service is based on the upstream [Manila](#) project. Red Hat provides articles on deploying and testing the OpenStack File Share Service, based on the following scenarios:

- [Using the glusterfs_native driver](#)

- [Using the glusterfs driver, with an NFS back end](#)
- [Using the glusterfs driver, with an NFS-Ganesha back end](#)

Each article also contains a list of known issues relevant to the driver used for each scenario.

Firewall-as-a-Service

The Firewall-as-a-Service (FWaaS) 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

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

Time Series Database-as-a-Service

Time Series Database-as-a-Service (based on the upstream [Gnocchi](#) project) allows you to define, store, and aggregate metrics in a more robust manner. It addresses many of the scalability and performance issues that arise when attempting to use Telemetry alone for performing monitoring, billing, and alarming in large enterprise environments.

VPN-as-a-Service

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

CHAPTER 3. RELEASE INFORMATION

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

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

https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux_OpenStack_Platform

3.1. ENHANCEMENTS

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

BZ#1261100

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#1041068

You can now use VMWare vSAN data stores. These stores allow you to use vMotion while simultaneously using hypervisor-local storage for instances.

BZ#1042222

The Orchestration service now includes an "OS::Heat::Stack" resource type. This OpenStack-native resource is used to explicitly create a child stack in a template. The "OS::Heat::Stack" resource type includes a 'context' property with a 'region_name' subproperty, allowing Orchestration service to manage stacks in different regions.

BZ#1052804

You can now use VMware storage policy to manage how storage is assigned to different instances. This can help you ensure that instances are assigned to the most appropriate storage in an environment where multiple data stores (of varying costs and performance properties) are attached to a VMware infrastructure.

BZ#1053078

Resources of type AWS::EC2::SecurityGroup can now be updated in-place when their rules are modified. This is consistent with the behaviour of AWS::EC2::SecurityGroup in CloudFormation. Previously, security groups would be replaced if they were modified.

BZ#1089447

This enhancement adds support for configuring multiple IPv6 prefixes and addresses on a single interface.

As a result, OpenStack Networking (neutron) considers the type of IPv6 subnets that form part of the network, and automatically associates ports with addresses from all the SLAAC-enabled subnets within the ports network.

There is no change to the REST API, but port-create/port-update responses automatically include the SLAAC addresses in the list of 'fixed_ips'.

BZ#1097987

Compute can now provide dedicated CPU resources, where each guest virtual CPU has full access to a specific host CPU. Previous releases of Compute guest CPUs were permitted to float across any host CPU. Even when the NUMA feature was enabled, the CPUs could still float within a NUMA node. Host CPUs would also overcommit so many virtual CPUs contended for the host resource. This made it impossible to provide strong performance guarantees to guest operating system workloads.

With this update, the cloud administrator now has the ability to set up a host aggregate, which provides a pool of hosts that supports guests with dedicated CPU resource assignment. The cloud administrator or tenant user can make use of these pools to run instances with guaranteed CPU resource.

BZ#1101375

OpenStack Trove instances can now be resized in the OpenStack dashboard user interface by selecting a new flavor for the instance.

BZ#1107490

The 'API Access' page in the dashboard ('Project > Compute > Access & Security > API Access') now provides more information on user credentials. To view this information, click 'View Credentials'. A pop-up displays the user name, project name, project ID, authentication URL, S3 URL, EC2 URL, EC2 access, and secret key.

BZ#1107924

The option to create Block Storage (cinder) volume transfers has been added to the 'Volumes' tab in the OpenStack dashboard. Volume transfers move ownership from one project to another. A donor creates a volume transfer, captures the resulting transfer ID and secret authentication key, and passes that information out of band to the recipient (such as by email or text message). The recipient accepts the transfer, supplying the transfer ID and authentication key. The ownership of the volume is then transferred from the donor to the recipient, and the volume is no longer visible to the donor.

Note the following limitations of the Block Storage API for volume transfers and their impact on the UI design:

1. When creating a volume transfer, you cannot specify who the intended recipient will be, and anyone with the transfer ID and authentication key can claim the volume. Therefore, the dashboard UI does not prompt for a recipient.
2. Current volume transfers are only visible to the donor; users in other projects are unable to view these transfers. So, the UI does not include a project table to view and accept volume transfers, since the current transfers are not visible. Instead, the transfer information is added to the volume details, which are visible by the donor, and the volume state clearly reflects that a transfer has been created. The UI also cannot present to the recipient a pull-down list of transfers to accept.
3. The only time that the authorization key is visible to the donor is in the response from the creation of the transfer; after creation, it is impossible for even the donor to recover it. Since the donor must capture the transfer ID and authorization key in order to send it to the recipient, an extra form was created to present this information to the donor immediately after the transfer has been created.

BZ#1108981

Heat now supports user hooks, which pause execution of stack operations at specified points to allow the user to insert their own actions into Heat's workflow. Hooks are attached to resources in the stack's environment file. Currently supported hook types are 'pre-create' and 'pre-update'.

BZ#1110589

The Identity Service (keystone) now allows for re-delegation of trusts. This allows a trustee with a trust token to create another trust to delegate their roles to others. In addition, a counter enumerates the number of times a trust can be re-delegated. This feature allows a trustee to re-delegate the roles contained in its trust token to another trustee. The user creating the initial trust can control if a trust can be re-delegated when this is necessary. Consequently, trusts can now be re-delegated if the original trust allows it.

BZ#1112481

OpenStack Dashboard now uses Block Storage (cinder) version 2 as its preferred version. Now when a Block Storage client is requested, access is given using cinder version 2, if not specified otherwise.

BZ#1114804

You can now use the dashboard to view, import, and associate metadata definitions that can be used with various resource types (images, artifacts, volumes, flavors, aggregates, etc).

BZ#1118578

The Image Service now features improved logging, providing better information to users. In addition, logs have been stripped of any sensitive information, and use the appropriate logging levels for messages. This change is only visible to operators.

BZ#1121844

Identity Service (keystone) now allows for unscoped tokens to be explicitly requested. This feature was added after users who had a default project assigned were previously unable to retrieve unscoped tokens; if one of these users requested a token without defining a scope, it would be automatically scoped to the default project. As a result of this update, unscoped tokens can now be issued to all users, even if they have a default project defined.

BZ#1121848

In OpenStack Dashboard, the instance detail page now displays the host node. This data is intended to assist when diagnosing issues.

BZ#1122774

The OS::Nova::Server resource type now includes a 'console_urls' property. This enables the user to obtain the URL for the server's console (such as a VNC console) from the resource.

BZ#1124672

This update adds partial support for Domain Admins to the OpenStack Dashboard. In addition, when using Identity Service (keystone) version 3, a newly-created user does not need to have a primary project specified.

BZ#1129773

With this enhancement, parameters CONFIG_CONTROLLER_HOST, CONFIG_COMPUTE_HOSTS, CONFIG_NETWORK_HOSTS support the use of hostname values along with the IP address values.

BZ#1133175

This update adds extended volume manage and unmanage support for NetApp Cmode and 7mode iSCSI drivers. This provides new functionality when using these drivers.

BZ#1133177

With this update, a new feature implements support to manage/unmanage volumes for the NetApp e-series driver. You can now use the '--source-name' parameter as the mandatory input for volumes not under the Block

Storage management.

BZ#1142563

When querying a resource in the Orchestration API, a user can now request the value of one or more of the resource's attributes be included in the output. This can aid debugging, as it allows the user to retrieve data from any resource at any time without having to modify the stack's template to include that data in the outputs section.

BZ#1143805

The `OS::Cinder::Volume` resource type now includes a `'scheduler_hints'` property. This allows scheduler hints to be passed to the Block Storage service when creating a volume, and requires v2 of the Block Storage API.

BZ#1143807

You can now disable and enable compute hosts through the dashboard. This capability is available through the 'Actions' column of every compute host in 'Admin > Hypervisors > Compute Host'.

Disabling a compute host prevents the scheduler from launching instances using that host.

BZ#1144230

The `heat-manage` command now includes a subcommand `"heat-manage service-list"`. This subcommand displays information about active "heat-engine" processes, where they are running, and their current status.

BZ#1149055

This enhancement adds namenode high availability as a supported option in the HDP 2.0.6 plugin. Users can signal that they require a cluster to be generated in HA mode, by passing a cluster with a quorum of zookeeper servers and journalnodes, and at least 2 namenodes. For example:

```
"cluster_configs": {
  "HDFSHA": {
    "hdfs.nnha": true
  }
}
```

BZ#1149959

The `OS::Neutron::Port` resource type now supports a `'binding:vnic_type'` property. This property enables users with the appropriate permissions to specify the VNIC type of an OpenStack Networking port.

BZ#1150839

The 'Manage/Unmanage' option has been added to the 'Volumes' tab of the OpenStack dashboard. 'Manage' takes an existing volume created outside of OpenStack and makes it available. 'Unmanage' removes the visibility of a volume within OpenStack, but does not delete the actual volume.

BZ#1151300

With this update, it is now possible to dynamically reload the Image service configuration settings by sending a SIGHUP signal to the 'glance-*' process. This signal will ensure the process re-reads the configuration file and load any new configurations. As a result, there is no need to restart the entire Image service to apply the configuration changes.

BZ#1151691

Bare Metal now supports the management interface of HP ProLiant Services using the iLO client python library. This allows Bare Metal to perform management operations such as retrieving/setting a boot device.

BZ#1153446

With this update, administrators are now able to view the state of High Availability routers on each node, and specifically, where the active instance is hosted. Previously, the High Availability router state information was not visible to the administrator; this made maintenance harder, for example, when moving HA router instances from one agent to another, or assessing the impact of putting a node in maintenance mode. This new functionality also serves as a sanity test and offers assurance that a router is indeed active on only one node. As a result, administrators may now run the 'neutron l3-agent-list-hosting-router <router_id>' command on a High Availability router to view where the active instance is currently hosted.

BZ#1153875

The Bare Metal service can now use cloud-init and similar early-initialization tools to insert user data on instances. Previously, doing so would have required setting up a metadata service to perform this function.

With this new update, Bare Metal can insert instance metadata onto local disk upon deployment -- specifically, to a device labeled 'config-2'. Afterwards, you can configure the early-initialization tool to find this device and extract the data from there.

BZ#1154485

The Bare Metal service can now deploy nodes using the Secure Boot feature of the UEFI (<http://www.uefi.org>). Secure Boot helps ensure that nodes boot only trusted software.

With this, the whole boot chain can be verified at boot time. You can then configure nodes to only boot authorized images, thereby enhancing security.

BZ#1154927

Bare Metal instances now feature a new field named 'maintenance_reason', which can be used to indicate why a node is in maintenance mode.

BZ#1155241

This package allows users to create HDP 2.0.6 and CDH 5.3.0 images for use in RHEL OpenStack Platform 7.

BZ#1155378

With this enhancement, the Sahara API now fully supports the HTTPS protocol.

BZ#1155388

With this update, the underlying asynchronous task engine has been changed. It is now based on the taskflow library. While this does not introduce changes to the API or workflow, it adds the following new configuration option:

```
[taskflow_executor]  
engine_mode = serial # or parallel
```

BZ#1156671

The `AWS::AutoScaling::AutoScalingGroup` resource type now supports an 'InstanceId' property. This allows the launch configuration for an autoscaling group to be cloned from an existing server instead of an `AWS::AutoScaling::LaunchConfiguration` resource.

BZ#1156678

The user interface options available in the dashboard for the OpenStack Orchestration service (heat) have been improved. For example, users can now check, suspend, resume, and preview stacks.

BZ#1156682

This update adds NFS back-ends for the cinder-backup service. This now allows back up of volumes to an NFS storage back end.

BZ#1158729

OpenStack Networking deployments with distributed routers are now able to allow tenants to create their own networks with VLAN segmentation.

Previously, distributed routers only supported tunnel networks, which may have hindered adoption as many deployments prefer to use VLAN tenant networks.

As a result of this update, distributed routers are now able to service tunnel networks as well as VLAN networks.

BZ#1159142

This update adds functionality to 'cinder-manage db' to safely purge old "deleted" data from the Cinder database. This reduces database space usage and improves database performance.

BZ#1159598

The `AWS::AutoScaling::LaunchConfiguration` resource type now supports an 'InstanceId' property. This allows the launch configuration for an autoscaling group to be cloned from an existing server.

BZ#1162436

The results displayed in tables for the DataProcessing service can now be filtered to allow the user to see only those results that are relevant.

BZ#1162961

You can now flag a volume as 'Bootable' through the dashboard.

BZ#1164087

Sahara objects can now be queried by any field name. This is done using the GET parameters that match the API field names, as seen on list methods.

BZ#1164520

Previously, the glance-manage utility was configured using 'glance-api.conf' or 'glance-registry.conf'. This release features a new configuration file named 'glance-manage.conf', which can be used to configure glance-manage. You can still use 'glance-api.conf' and 'glance-registry.conf' to configure glance-manage, but any 'glance-manage.conf' settings will take precedence.

BZ#1165499

The Bare Metal service now supports Fujitsu iRMC (integrated Remote Management Controller) hardware. With this, Bare Metal can now manage the power state of such machines.

BZ#1165505

With this update, Identity Service (keystone) is now able to construct

With this update, Identity service (keystone), is now able to construct a hierarchy of projects by specifying a 'parent_id' within a project resource.

Previously, the Identity service only allowed for a flat project model; a project hierarchy allows for more flexible project structures, which can be used to mimic organizational structures.

As a result, Projects can now define a parent project, allowing project hierarchies to be constructed.

BZ#1166490

The OpenStack dashboard can now use a custom theme. A new setting, 'CUSTOM_THEME_PATH' was added to /etc/openstack_dashboard/local_settings file. The theme folder should contain one _variables.scss file and one _styles.scss file. The _variables.scss file contains all the bootstrap and Horizon-specific variables that are used to style the graphical user interface, and the _styles.scss file contains extra styling.

BZ#1168371

Previously, Image service's 'swift' store implementation stored all images on a single container. While this worked well, it created a performance bottleneck in large scale deployments.

With this update, it is now possible to use several Object Storage containers as storage for the 'glance' images. In order to use this feature, you need to set 'swift_store_multiple_containers_seed' to a value bigger than '0'. You can disable using multiple containers by enabling the 'swift_uer_multi_tenant' parameter, as these containers are split on a per-tenant basis.

BZ#1170470

SRIOV can now be configured in the OpenStack dashboard. Options include exposing further information on the 'Port Details' tab, and allowing port type selection during port creation and update.

BZ#1170471

This enhancement allows you to view encryption metadata for encrypted volumes in OpenStack Dashboard (horizon). A function to display encryption metadata was added, and allows the user to click on the "Yes" in the Encrypted column, and be taken to a page where the encryption metadata is visible.

BZ#1170475

The glance_store library now supports more storage capabilities. As such, you now have more granular control over what operations are allowed in a specific store. This release features the following capabilities:

- READ_ACCESS: Generic read access
- WRITE_ACCESS: Generic write access

- RW_ACCESS : READ_ACCESS and WRITE_ACCESS
- READ_OFFSET: Read all bits from a offset (Included in READ_ACCESS)
- WRITE_OFFSET: Write all bits to a offset (Included in WRITE_ACCESS)
- RW_OFFSET : READ_OFFSET and WRITE_OFFSET
- READ_CHUNK : Read required length of bits (Included in READ_ACCESS)
- WRITE_CHUNK: Write required length of bits (Included in WRITE_ACCESS)
- RW_CHUNK: READ_CHUNK and WRITE_CHUNK
- READ_RANDOM: READ_OFFSET and READ_CHUNK
- WRITE_RANDOM: WRITE_OFFSET and WRITE_CHUNK
- RW_RANDOM: RW_OFFSET and RW_CHUNK
- DRIVER_REUSABLE: driver is stateless and its instance can be reused safely

BZ#1170476

With this update, a completely new API that adds search capabilities for Image service and improves the performance for listing and search operations, especially on interactions with the UI is now available.

The search API allows users to execute a search query and get back search hits that match the query. The query can either be provided using a simple query string as a parameter, or using a request body. All the search APIs can be applied across multiple types within an index, and across multiple indices with support for multi index syntax.

Note: This enhancement will be removed from the Image service during the RHEL OpenStack Platform 8 (Liberty) release.

BZ#1185652

This feature adds IPv6 support to Packstack, allowing Packstack to use IPv6 address as values in networking-related parameters such as CONFIG_CONTROLLER_HOST, CONFIG_COMPUTE_HOSTS, and CONFIG_NETWORK_HOSTS.

BZ#1189500

This enhancement adds a CLI that allows configuration of the default cluster templates for each major plugin. The provision of default templates is expected to speed and facilitate end-user adoption of Sahara.

As a result of this update, administrators can now add shared default templates for adaptation and direct usage by customers.

BZ#1189504

Integration tests for Sahara have been refactored from more brittle pure python tests to allow easy, YAML-based configuration to define "scenarios".

BZ#1189511

Previously, the cm_api library was not packaged by Cloudera for any

Linux distribution. The previous CDH plug-in depended on this package, so CDH could not be enabled as a default plug-in prior to this release. Now, a subset of the `cm_api` library has been added to Sahara's codebase, and CDH is functional and enabled by default.

BZ#1189633

The Identity service now allows unscoped federation tokens to be used to obtain a scoped token using the 'token' authentication method.

When using the Identity service's federation extension, an unscoped federation token is returned as a result of the initial authentication. This is then exchanged for a scoped token. An unscoped federation token previously had to use the 'saml2' or 'mapped' authentication to obtain a scoped token. This is inconsistent with the method used to exchanging a regular unscoped token for a scoped token, which uses the 'token' method.

Exchanging an unscoped federation token for a scoped token now uses the 'token' authentication method, which is consistent with the regular unscoped token behavior.

BZ#1189639

The Identity service now allows restriction of re-scoping tokens to only allow unscoped changes to be exchanged for scoped tokens.

The Identity service allows for an existing token to be used to obtain a new token via the 'token' authentication method. Previously, a user with a valid token scoped for a project could use that token to obtain another token for a different project that they were authorized for. This allowed for anyone possessing a user's token to have access to any project the user has access to, as opposed to only having access to the project that the token is scoped for. To improve the security properties of scoped tokens, it was desirable to not allow this.

A new 'allow_rescope_scoped_token' configuration option is available to allow token rescoping to be restricted. Rescoping of tokens is now only allowed by using an unscoped token to authenticate when this option is enabled.

BZ#1189711

The dashboard now provides wizards for creating and configuring the necessary components of the OpenStack Data Processing feature. These wizards are useful for guiding users through the process of cluster creation and job execution. To use these wizards, go to 'Project > Data Processing > Guides'.

BZ#1189716

This enhancement adds ceilometer IPMI meters to OpenStack Dashboard. Six ipmi meters have been exported from ceilometer; the methods 'list_ipmi' and '_get_ipmi_meters_info' are used to retrieve the meter

data.

BZ#1189811

Previously, every call to `policy.enforce` passed an empty dictionary as the target. This prevented operators from using tenant specific restrictions in their `policy.json` files since the target would always be an empty dictionary. If you tried to restrict some actions so an image owner (users with the correct tenant id) could perform actions, the check categorically failed because the target is okay is an empty dictionary.

With this update, you can pass the `ImageTarget` instance wrapping an `Image` to the enforcer so these rules can be used and properly enforced. You can now properly grant access to the image owner(s) based on tenant (e.g., `owner:%(tenant)`). Without this fix, the only check that actually works in Image service is a `RoleCheck` (e.g., `role:admin`).

BZ#1190312

You can now view details about Orchestration service hosts through the dashboard. To do so, go to 'Admin > System > System Information > Orchestration Services'. This page is only available if the Orchestration service is deployed.

BZ#1192290

Previously, many of the processes in cluster creation polled infinitely. Now, timeouts have been added for many stages of cluster creation and manipulation, and users are shown appropriate error messages when cluster operations have taken longer than is reasonable.

BZ#1193287

Support has been added for intelligent NUMA node placement for guests that have been assigned a host PCI device. PCI I/O devices, such as Network Interface Cards (NICs), can be more closely associated with one processor than another. This is important because there are different memory performance and latency characteristics when accessing memory directly attached to one processor than when accessing memory directly attached to another processor in the same server. With this update, Openstack guest placement can be optimized by ensuring that a guest bound to a PCI device is scheduled to run on a NUMA node that is associated with the guest's pCPU and memory allocation. For example, if a guest's resource requirements fit in a single NUMA node, all guest resources will now be associated with the same NUMA node.

BZ#1194532

A new endpoint has been added to Sahara that allows queries of the available job types per plug-in and version that the Sahara installation supports. This information is useful both for UI presentation and filtering, and for CLI and REST API users.

BZ#1196013

The Identity service now has an experimental support for a new token format called 'fernet'.

The token formats currently supported by the Identity service require issued tokens to be persisted in a database table. This table can grow quite large, which requires proper tuning and a flush job to keep the Identity service performing well. The new 'fernet' token format is designed to allow the token database table to be eliminated, avoiding the problem of this table becoming a scalability limitation. The 'fernet' token format is now available as an experimental feature.

BZ#1198904

All Ironic drivers now support deployment via IPA ramdisk. IPA is written in Python, supports more features than the BASH ramdisk, and runs as a service. For these reasons, nodes deployed through IPA are generally easier to deploy, debug, and manage.

BZ#1198911

With this update, it is now possible to filter the list operations by more than one filter option and in multiple directions. For example:

```
/images?sort=status:asc,name:asc,created_at:desc
```

With the above, a list of images will be returned and they will be sorted by status, name, and creation date with the following directions respectively: ascending, ascending, and descending.

BZ#1201116

With this change, it is now possible to filter the list operations by more than one filter option and in multiple directions. For example:

```
/images?sort=status:asc,name:asc,created_at:desc
```

With the above, a list of images will be returned and they will be sorted by status, name, and creation date with the following directions respectively: ascending, ascending, and descending.

BZ#1202472

This update adds the ability to assign a user group as the instance owner, which allows other members of the same group to control the instance when its creator is not reachable.

BZ#1205869

Imbalance in tiers in Swift was previously addressed by weights. However, no matter what the numerical ratio of weight is set, at a certain point there are not enough devices and replicas remaining in

lower weight tiers to balance out the crowding at the higher weight tier. At this moment, the tier becomes underutilized, while an administrator may need to force more than one replica into tier to achieve utilization. The ratio of more-than-1 partition is the overload parameter.

This update permits administrators to store more than one replica in a tier in case of severely unbalanced clusters. As a result, it is now possible so sacrifice data durability in order to achieve better utilization, which in some cases is required for availability. For example, a cluster will fail to store new data if low-weight tiers overflow and quorum fails.

BZ#1209908

Previously, the GUI did not allow a user to create or upload images.

This new feature adds the following changes to the horizon and the tuskar-ui packages. To horizon, it adds the Kernel and Ramdisk fields to the create image form in horizon, which enables a user to associate a kernel and a ramdisk to a glance image during image creation. In the tuskar-ui, it exposes horizon's create image form (with the newly added capability to set kernel and ramdisk) in the tuskar-ui. On the images page, there should now be a "create image" button which allows the user to create an image. As a result, GUI now has the ability to create images.

BZ#1215790

Previously, when using huge pages, the back-end memory for a guest was configured as private. Consequently, the vhostuser VIF back end was designed to allow an external process to provide the QEMU network driver functionality. For some use cases of vhostuser, this required that the external process be able to access the QEMU guest's memory pages directly. This is not possible when the huge pages are mapped with MAP_PRIVATE; they must use MAP_SHARED instead. With this update, when a guest is configured to use huge pages backed memory, mappings are be marked as shared. As a result, the external process to provide QEMU network is now able to access to guest's memory pages.

BZ#1229811

This enhancement adds support for the Cisco N1kV plugin. This includes environment configuration in the TripleO Heat Template collection.

BZ#1230844

This enhancement adds support for the Nexus-9k ML2 Neutron plugin. This includes environment configuration in the TripleO Heat Template collection as well as configuration in the Openstack Puppet Module collection.

BZ#1230850

This enhancement adds support for the Cisco UCSM Neutron ML2 plugin. This includes environment configuration in the TripleO Heat Template collection as well as configuration in the Openstack Puppet Module collection.

BZ#1230875

Bare Metal Provisioning (Ironic) now supports a driver that manages Cisco UCS servers. Using the new driver with Cisco UCS servers allows for better support for more advanced features.

BZ#1233564

This fix adds support for Cisco UCS machines to Ironic's power management control in the director. Cisco UCS nodes are manageable using the IPMI protocol, but some customers might want to use the specific Cisco UCS driver to manage more advanced features. Now the director supports power management for Cisco UCS machines.

BZ#1236055

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#1238740

The nexus1000v (n1kv) Puppet class has been added.

BZ#1241094

Users can now set the maintenance mode and provision state of Bare Metal Provisioning (Ironic) nodes using 'openstackclient' commands. Previously, Ironic used a mix of 'python-ironicclient' and 'openstackclient' commands. This enhancement provides a more unified interface to the user. The new commands are available as part of the 'openstack baremetal' command-line interface.

BZ#1241720

This enhancement adds support for the Cisco N1kV VEM module. This includes environment configuration in the TripleO Heat Template collection.

BZ#1244010

This enhancement adds Linux bonding configuration through the director. The director used only OVS and VLAN bonding previously. Linux bonding provides increased performance and additional bonding modes.

BZ#1247982

The kafka-python library is now included in this release (provided by the python-kafka package). This library provides support for Apache Kafka; this, in turn, allows the Telemetry service to dispatch events and samples with the Kafka publisher.

BZ#1249832

This enhancement increases the levels of configuration for the Overcloud's Neutron service. Customers can now configure values for `core_plugin`, `type_drivers`, and `service_plugins` through the director.

BZ#1254153

The enhancement adds the 'python-networking-cisco' package. This enables support for multiple Cisco plugins and drivers in OpenStack Networking (neutron).

BZ#1257606

This feature allows the S3 driver to be configured to pass through a proxy. The boto library already supported this capability, but it was not exposed through the `glance_store` API.

The following configuration options have been added and are turned off by default. They must be configured for the S3 driver to use the proxy:

- * `s3_store_enable_proxy`
- * `s3_store_proxy_host`
- * `s3_store_proxy_port`
- * `s3_store_proxy_user`
- * `s3_store_proxy_password`

BZ#1257717

On a PATCH update (using the "-x" flag in the 'heat stack-update' command), the existing environment is now retained unless explicitly overridden. This is because the Orchestration service now re-uses other parts of the environment, not just the parameters that were passed previously and not overridden.

This feature was added because in the most common stack update cases, users prefer to maintain the current environment (including resource mappings and the like). This will also prevent any unintended changes in complex deployments whenever users forget to include the required environment files at stack creation time.

BZ#1259393

This enhancement adds support for the `fake_pxe` Ironic driver for registering machines without power management to the director. Use the `fake_pxe` driver as a fallback driver for machines without a power

management system. Perform all power operations manually when using this driver.

BZ#1272176

This enhancement upgrades the Overcloud image content to Red Hat Enterprise Linux 7.2 content, including the latest version of Pacemaker. The previous Overcloud image used Red Hat Enterprise Linux 7.1 content.

BZ#1274241

This enhancement adds support for Fujitsu's iRMC Ironic driver in the director. The director now controls the power management of iRMC nodes in the Overcloud.

BZ#1274444

The Overcloud image is now multipath aware. This helps users aiming to deploy on nodes using a multipathed boot LUN. The operating system root is now mounted properly (e.g. /dev/mpatha).

BZ#1275439

This feature allows the reapplication of Puppet manifests on a deployed Overcloud. This ensures the overcloud has the desired configuration or can recover accidentally amended or deleted configuration files.

To have Puppet run again on the Overcloud nodes, omit the "--templates" option but include the following two environments files at the beginning of your deployment:

```
* /usr/share/openstack-tripleo-heat-templates/overcloud-resource-registry-puppet.yaml
* /usr/share/openstack-tripleo-heat-templates/extraconfig/pre_deploy/rhel-registration/rhel-registration-resource-registry.yaml
```

For example:

```
$ openstack overcloud deploy -e ~/templates/overcloud-resource-registry-puppet.yaml -e ~/templates/extraconfig/pre_deploy/rhel-registration/rhel-registration-resource-registry.yaml [additional arguments from initial deployment]
```

BZ#1278868

This enhancement adds support for the Nuage on highly available Overcloud environments. This includes Nuage-specific parameters in the director's Heat template collection, and environment files to enable the Nuage backend on Controller and Compute nodes.

BZ#1278879
■

This enhancement adds support for the Nuage metadata agent on the Overcloud. This includes parameters in the director's Heat template collection for the Nuage metadata agent.

BZ#1293473

This enhancement adds support to register Overcloud nodes to a Red Hat Satellite 5 server. Previous versions allowed registration only to a Red Hat Satellite 6 server. Now the director determines whether to register to a Red Hat Satellite 5 or Red Hat Satellite 6 server when using the '-reg-method satellite' option during Overcloud creation.

BZ#1298197

This enhancement adds SSL support to the Overcloud's Public API. Users can now configure SSL on the Overcloud using the 'environments/enable-tls.yaml' from the director's Heat template collection. Copy and modify this environment file to suit your SSL requirements. For more information, see "6.2.7. Enabling SSL/TLS on the Overcloud" in the Director Installation and Usage guide for Red Hat OpenStack Platform 7.3.

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#1252504

With this update, the package python-pytimemachine is rebased to version 1.1.5 for openstack-gnocchi, so that gnocchi-dbsync can work successfully without any ImportError.

BZ#1259740

OpenStack Compute includes the concept of Cells, provided by the nova-cells package, for dividing computing resources.

Cells are provided in Red Hat Enterprise Linux OpenStack Platform as a Technology Preview at this time. Fully supported methods for dividing compute resources in Red Hat Enterprise Linux OpenStack Platform include Regions, Availability Zones, and Host Aggregates.

3.3. RELEASE NOTES

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

BZ#894888

Support for SPICE remote console access was recently added to the Compute (Nova) and Dashboard (Horizon) services. The `spice-html5` package required to support SPICE access is however not included in this release. As such SPICE remote console access remains unsupported at this time.

BZ#1057941

This feature aims to improve the `libvirt` driver so that it can use large pages for backing the guest RAM allocation. This will improve the performance of guest workloads by increasing TLB cache efficiency. It will ensure that the guest has 100% dedicated RAM that will never be swapped out.

BZ#1108439

Efficient replication of distributed regions is now possible. Existing replication sends data over unnecessarily. For example, in case of two regions and replication factor of four, a typical spread is to have two replicas per region, assuming uniform weights. If the long-haul link goes down and up, replicators send two copies of every outstanding object across regions. This is not necessary. If only one copy is sent over, replicators of the remote region restore complete redundancy without using the long-haul bandwidth.

This update ensures stable performance during a recovery after an outage in a multi-region cluster. It has no effect inside a one-region datacenter.

BZ#1203160

After fully upgrading to Red Hat Enterprise Linux OpenStack Platform 7 from version 6 (and all nodes are running version 7 code), you should start a background migration of PCI device NUMA node information from the old location to the new location. Version 7 conductor nodes will do this automatically when necessary, but the rest of the idle data needs to be migrated in the background. This is critical to complete before the version 8 release, where support for the old location will be dropped. Use `'$ nova-manage db migrate_rhos_6_pci_device_data --max-number X'` to perform this transition (where X is the maximum number of devices to be migrated in one run).

Note that this is relevant only for users making use of the PCI pass-through features of Compute.

BZ#1228096

In Kilo, Neutron services now can rely on so called `rootwrap` daemon to execute external commands like `'ip'` or `'sysctl'`. The daemon pre-caches `rootwrap` filters and drastically improves overall agent performance.

For RHEL-OSP7, `rootwrap` daemon is enabled by default. If you want to

avoid using it and stick to another root privilege separation mechanism like 'sudo', then make sure you also disable the daemon by setting 'root_helper_daemon =' in [agent] section of your neutron.conf file.

BZ#1265777

When using static IP's on the ctlpane network, DNS must be configured on the Overcloud nodes. Previously, DHCP provided the configured DNS servers for the Overcloud nodes.

To configure DNS when using static IPs set the new DnsServers parameter and include it in the Heat environment like so:

```
parameter_defaults:
  DnsServers:
    - <dns server ip address>
    - <dns server ip address 2>
```

Use IP addresses for the DNS servers to use. You can only specify either one or two DNS servers.

3.4. KNOWN ISSUES

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

BZ#1221034

Due to a known issue with the 'python-neutron-fwaas' package, Firewall-as-a-Service (FWaaS) may fail to work. This is a result of the 'python-neutron-fwaas' package missing the database upgrade 'versions' directory.

In addition, upgrading the database schemas between version releases may not function correctly at this time.

BZ#1244358

The Director uses misconfigured HAProxy settings when deploying the Bare Metal and Telemetry services with SSL enabled in the undercloud. This prevents some nodes from registering.

To work around this, comment out 'option ssl-hello-chk' under the Bare Metal and Telemetry sections in /etc/haproxy/haproxy.cfg after installing the undercloud.

BZ#1256630

The GlusterFS native driver of the File Share Service allows users to create shares of specified sizes. If no Red Hat Gluster volumes of the exact requested size exist, the driver chooses one with the nearest possible size and creates a share on the volume. Whenever this occurs, the resulting share will use the entire volume.

For example, if a user requests a 1GB share and only 2GB, 3GB, and 4GB

volumes are available, the driver will choose the 2GB volume as a back end for the share. The driver will also proceed with creating a 2GB share; the user will be able to use and mount the entire 2GB share.

To work around this, implement File Share quotas for users. Doing so will prevent them from provisioning more file share storage than what they are entitled to.

BZ#1272347

With this update, the default network where the 'KeystoneAdminVip' is placed was changed from 'InternalApi' to 'ctlplane' so that the post-deployment Identity service initialization step could be carried by the Undercloud over the 'ctlplane' network. Relocating the 'KeystoneAdminVip' causes a cascading restart of the services pointing to the old 'KeystoneAdminVip'.

As a workaround to make sure the KeystoneAdminVip remains on the 'InternalApi' network, a customized 'ServiceNetMap' must be provided as deployment parameter when launching an update from the 7.0 release. A sample Orchestration environment file passing a customized 'ServiceNetMap' is as follows:

parameters:

ServiceNetMap:

```
NeutronTenantNetwork: tenant
CeilometerApiNetwork: internal_api
MongoDbNetwork: internal_api
CinderApiNetwork: internal_api
CinderIscsiNetwork: storage
GlanceApiNetwork: storage
GlanceRegistryNetwork: internal_api
KeystoneAdminApiNetwork: internal_api
KeystonePublicApiNetwork: internal_api
NeutronApiNetwork: internal_api
HeatApiNetwork: internal_api
NovaApiNetwork: internal_api
NovaMetadataNetwork: internal_api
NovaVncProxyNetwork: internal_api
SwiftMgmtNetwork: storage_mgmt
SwiftProxyNetwork: storage
HorizonNetwork: internal_api
MemcachedNetwork: internal_api
RabbitMqNetwork: internal_api
RedisNetwork: internal_api
MysqlNetwork: internal_api
CephClusterNetwork: storage_mgmt
CephPublicNetwork: storage
ControllerHostnameResolveNetwork: internal_api
ComputeHostnameResolveNetwork: internal_api
BlockStorageHostnameResolveNetwork: internal_api
ObjectStorageHostnameResolveNetwork: internal_api
CephStorageHostnameResolveNetwork: storage
```

If any additional binding network from the above has been customized

then that setting has to be preserved as well.

As a result of the workaround changes, the 'KeystoneAdminVip' is not relocated on the 'ctlplane' network so that no services restart needs to be triggered.

BZ#1290949

By default the number of heat-engine workers created will match the number of cores on the undercloud. Previously, however, if there was only one core there would only be one heat-engine worker, and this caused deadlocks when creating the overcloud stack. A single heat-engine worker was not enough to launch an overcloud stack.

To avoid this, it is recommended that the undercloud has at least two (virtual) cores. For virtual deployments this should be two vCPUs, regardless of cores on the baremetal host. If this is not possible, then uncommenting the `num_engine_workers` line in `/etc/heat/heat.conf`, and restarting `openstack-heat-engine` fixes the issue. Thus, the above workarounds have resolved the issue.

BZ#1069157

At present, policy rules for volume extension prevent you from taking snapshots of in-use GlusterFS volumes. To work around this, you will have to manually edit those policy rules.

To do so, open the Compute service's `policy.json` file and change "rule:admin_api" entries to "" for "compute_extension:os-assisted-volume-snapshots:create" and "compute_extension:os-assisted-volume-snapshots:delete". Afterwards, restart the Compute API service.

BZ#1220630

The underlying Database-as-a-Service (Trove) processes will not start if the service's back-end database is unreachable. To work around this, Database-as-a-Service must be deployed on the same node as its back-end database.

BZ#1241424

Sometimes bare metal nodes can lock into a certain state if ironic-conductor stops abruptly. This means users cannot delete these nodes or change their state. As a workaround, log into the director's database and use the following query to set the node back to "available" state and remove the lock:

```
UPDATE nodes SET provision_state="available",
target_provision_state=NULL, reservation=NULL WHERE uuid=<node uuid>;
```

BZ#1221076

Due to a known issue with the 'python-neutron-fwaas' package, Firewall-

as-a-Service (FWaaS) may fail to work. This is a result of the 'python-neutron-fwaas' package missing the database upgrade 'versions' directory.

In addition, upgrading the database schemas between version releases may not function correctly at this time.

BZ#1247358

In rare cases, RabbitMQ fails to start on deployment. As a workaround, manually start RabbitMQ on nodes:

```
[stack@director ~]$ ssh heat-admin@192.168.0.20
[heat-admin@overcloud-controller-0 ~]$ pcs resource debug-start rabbitmq
```

Then rerun the deployment command on the director. The deployment now succeeds.

BZ#1246525

On the Undercloud, HAProxy is configured to run a HTTP check against the openstack-ironic-api service every 2 seconds. The check causes openstack-ironic-api to log a traceback to stderr with the errors:

```
error: [Errno 104] Connection reset by peer
error: [Errno 32] Broken pipe
```

Since the check runs every 2 seconds, these messages repeat frequently in /var/log/messages. As a workaround, switch to root permissions, edit /etc/haproxy/haproxy.cfg, and comment out the "option httpchk GET /" line from the ironic listener configuration:

```
listen ironic
  bind 192.0.2.2:6385
  bind 192.0.2.3:6385
  # option httpchk GET /
  server 192.0.2.1 192.0.2.1:6385 check fall 5 inter 2000 rise 2
```

Save the file, then restart haproxy:

```
sudo systemctl restart haproxy
```

No tracebacks from openstack-ironic-api are written to stderr.

BZ#1296365

Multiple services attempted NTP configuration on the Overcloud and the last service configured it incorrectly. This caused time synchronization issues across all Overcloud nodes. As a workaround, delete /usr/libexec/os-apply-config/templates/etc/ntp.conf from all Overcloud nodes and re-run the deployment command to re-apply the puppet configuration. This is required for users updating from an older version of Red Hat OpenStack Platform to 7.3. This fix is not necessary on new 7.3 deployments. NTP now configures correctly.

BZ#1236136

All keystone endpoints are on the External VIP. This means all API calls to keystone happen over the External VIP. There is no workaround at this time.

BZ#1250043

When using the 'gluster_native' driver for File Share Service back ends, snapshot commands can fail ungracefully with a 'key error' if any of the following components are down in the back end cluster's nodes:

- Logical volume brick
- The glusterd service
- Red Hat Gluster Storage volume

In addition, the following could also cause the same error:

- An entire node in a cluster is down.
- An unsupported volume is used as a back end.

Specifically, these issues can cause the 'openstack-manila-share' service to produce a traceback with KeyError instead of producing a useful error message. When troubleshooting this error, consider these possible back end issues.

BZ#1205432

The OpenStack Dashboard (Horizon) is not configured to accept connections on its local IP address. This means you cannot browse the OpenStack Dashboard, including the Undercloud UI by IP address. As a workaround, use the Undercloud's FQDN instead of IP address. If access through the IP address is desired, edit /etc/openstack-dashboard/local_settings, add the IP address to the ALLOWED_HOSTS setting, then restart the httpd service. This enables the ability to browse OpenStack Dashboard through the host IP address.

BZ#1257291

With the glusterFS_native driver, providing or revoking 'cert'-based access to a share restarts a Red Hat Gluster Storage volume. This, in turn, will disrupt any ongoing I/O to existing mounts. To prevent any data loss, unmount a share on all clients before allowing or denying access to it.

BZ#1250130

The 'manila list' command shows information on all shares available from the File Share Service. This command also shows the Export Location field of each one, which should provide information for composing its mount point entry in an instance. However, the field displays this information in the following format:

```
user@host:/vol
```

The 'user@' prefix is unnecessary, and should therefore be ignored when composing its mount point entry.

BZ#1257304

With the the File Share Service, when an attempt to create a snapshot of a provisioned share fails, an entry for the snapshot will still be created. However, this entry will be in an 'error' state, and any attempts to delete it will fail.

To prevent this, avoid creating share snapshots if the back end volume, service, or host is down.

BZ#1307125

If the database is not available when heat-engine starts, heat-engine fails to start, preventing heat-engine from running under certain circumstances. This occurs after you update a machine on which both the database and heat-engine are running using yum alone such as when updating the undercloud. As a workaround, heat-engine must be explicitly started by running 'systemctl start openstack-heat-engine.service'. You can confirm whether heat-engine is running by running 'systemctl status openstack-heat-engine.service'.

BZ#1321179

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

CHAPTER 4. UPGRADING

For instructions on how to upgrade RHEL OpenStack Platform components from version 6 to version 7, see the following guide:

- [Upgrading OpenStack](#)

For instructions on how to update the undercloud and overcloud from the initial release of Red Hat Enterprise Linux OpenStack Platform 7.0 manually, see the [Updating the Environment](#) chapter in *Director Installation and Usage*.

CHAPTER 5. ADDITIONAL OVERCLOUD CONFIGURATION PARAMETERS

Version 7.0 A2 of the director includes support for static Provisioning IPs. The systems boot via DHCP during deployment, and the DHCP address assigned converts to a static IP. The following parameters have been added to support static IP addressing on the provisioning network:

- ControlPlaneIp
- ControlPlaneSubnetCidr
- DnsServers
- EC2MetadataIp

These changes require additional parameters for setting static IPs, routes, and DNS servers. When using static Provisioning IPs, the network environment file now needs to contain additional resource defaults, which you customize to match your environment. For example:

```
parameter_defaults:
  # CIDR subnet mask length for provisioning network
  ControlPlaneSubnetCidr: 24
  # Gateway router for the provisioning network (or Undercloud IP)
  ControlPlaneDefaultRoute:10.8.146.254
  # Generally the IP of the Undercloud
  EC2MetadataIp: 10.8.146.1
  # Define the DNS servers (maximum 2) for the overcloud nodes
  DnsServers:['8.8.8.8','8.8.4.4']
```

The NIC configuration templates for each role now include additional parameters in the **parameters** section. Regardless of whether the Provisioning interface uses DHCP or static IPs, these parameters are needed in any case:

```
parameters:
  ControlPlaneIp:
    default: ''
    description: IP address/subnet on the ctlplane network
    type: string
  ControlPlaneSubnetCidr: # Override this via parameter_defaults
    default: '24'
    description: The subnet CIDR of the control plane network.
    type: string
  DnsServers: # Override this via parameter_defaults
    default: []
    description: A list of DNS servers (2 max) to add to resolv.conf.
    type: json
  EC2MetadataIp: # Override this via parameter_defaults
    description: The IP address of the EC2 metadata server.
    type: string
```

If customizing the templates in the **network/config** subdirectory of the director's Heat template collection, note that these files are updated with these new parameters. If you have NIC configuration templates from an older version of the director's Heat templates collection, add these new parameters and modify the provisioning network to take advantage of static IP addresses.

For more information, see:

- [Section 6.3.7. Isolating all Networks into VLANs](#)
- [Appendix E. Network Interface Template Examples](#)

CHAPTER 6. TECHNICAL NOTES

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

6.1. RHEA-2015:1548 — RED HAT ENTERPRISE LINUX OPENSTACK PLATFORM ENHANCEMENT ADVISORY

The bugs contained in this section are addressed by advisory RHEA-2015:1548. Further information about this advisory is available at <https://access.redhat.com/errata/RHEA-2015:1548.html>.

6.1.1. crudini

BZ#1223624

Prior to this update, separate lock files were used while updating config files. In addition, directory entries were not correctly synchronized during an update.

As a result, a crash during this process could cause deadlock issues on subsequent config update attempts, or very occasionally result in corrupted (empty) config files.

This update adds more robust locking and synchronization within the 'crudini' utility. The result is that config file updates are now more robust during system crash events.

6.1.2. mariadb-galera

BZ#1211088

This rebase package includes a notable fix under version 5.5.42:

* An issue was resolved whereby INSERT statements that use auto-incrementing primary keys could fail with a "DUPLICATE PRIMARY KEY" error on an otherwise working Galera node, if a different Galera node that was also handling INSERT statements on that same table was recently taken out of the cluster. The issue would cause OpenStack applications to temporarily fail to create new records while a Galera failover operation was in-progress.

6.1.3. openstack-ceilometer

BZ#1232163

Previous versions of 'alarm-history' did not give an indication of when the severity of a given alarm was changed (for example, from 'low' to 'critical'); instead a change was indicated without any detail given of what the change was.

This update addresses this issue with a code update that displays severity changes.

As a result, 'alarm-history' now displays severity changes in the output of alarm-history.

BZ#1240532

Previously, when a ceilometer polling extension could not be loaded, an ERROR message was logged. This was misleading in cases where the failure to load a module was the expected outcome, such as when an extension was optional or its dependent modules were not available. Now, the log messages have been changed to WARN level to make it clear that there is no serious fault.

6.1.4. openstack-cinder**BZ#1133175**

This update adds extended volume manage and unmanage support for NetApp Cmode and 7mode iSCSI drivers. This provides new functionality when using these drivers.

BZ#1133177

With this update, a new feature implements support to manage/unmanage volumes for the NetApp e-series driver. You can now use the '--source-name' parameter as the mandatory input for volumes not under the Block Storage management.

BZ#1156682

This update adds NFS back-ends for the cinder-backup service. This now allows back up of volumes to an NFS storage back end.

BZ#1159142

This update adds functionality to 'cinder-manage db' to safely purge old "deleted" data from the Cinder database. This reduces database space usage and improves database performance.

BZ#1200986

Prior to this update, SQLAlchemy objects were incorrectly shared between multiple 'cinder-volume' processes.

Consequently, SQLAlchemy connections would fail when using a Block Storage multi-backend, resulting in database-related errors in the volume service.

This fix re-initializes SQLAlchemy connections when forking 'cinder-volume' child processes.

As a result multi-backend now works as expected.

BZ#1208767

In the previous version, creating volume from an image failed. On a virtual disk with a high number of sectors, the number of sectors was in

some cases handled incorrectly and, converting a QEMU image failed with an "invalid argument" error.

This bug has been resolved by updating to a fixed version of QEMU-img that resolves the incorrect calculation issue that caused this error. Creating volume from image now works successfully.

6.1.5. openstack-glance

BZ#1118578

The Image Service now features improved logging, providing better information to users. In addition, logs have been stripped of any sensitive information, and use the appropriate logging levels for messages. This change is only visible to operators.

BZ#1151300

With this update, it is now possible to dynamically reload the Image service configuration settings by sending a SIGHUP signal to the 'glance-*' process. This signal will ensure the process re-reads the configuration file and load any new configurations. As a result, there is no need to restart the entire Image service to apply the configuration changes.

BZ#1155388

With this update, the underlying asynchronous task engine has been changed. It is now based on the taskflow library. While this does not introduce changes to the API or workflow, it adds the following new configuration option:

```
[taskflow_executor]
engine_mode = serial # or parallel
```

BZ#1164520

Previously, the glance-manage utility was configured using 'glance-api.conf' or 'glance-registry.conf'. This release features a new configuration file named 'glance-manage.conf', which can be used to configure glance-manage. You can still use 'glance-api.conf' and 'glance-registry.conf' to configure glance-manage, but any 'glance-manage.conf' settings will take precedence.

BZ#1168371

Previously, Image service's 'swift' store implementation stored all images on a single container. While this worked well, it created a performance bottleneck in large scale deployments.

With this update, it is now possible to use several Object Storage

containers as storage for the 'glance' images. In order to use this feature, you need to set 'swift_store_multiple_containers_seed' to a value bigger than '0'. You can disable using multiple containers by enabling the 'swift_uer_multi_tenant' parameter, as these containers are split on a per-tenant basis.

BZ#1170475

The glance_store library now supports more storage capabilities. As such, you now have more granular control over what operations are allowed in a specific store. This release features the following capabilities:

- READ_ACCESS: Generic read access
- WRITE_ACCESS: Generic write access
- RW_ACCESS : READ_ACCESS and WRITE_ACCESS
- READ_OFFSET: Read all bits from a offset (Included in READ_ACCESS)
- WRITE_OFFSET: Write all bits to a offset (Included in WRITE_ACCESS)
- RW_OFFSET : READ_OFFSET and WRITE_OFFSET
- READ_CHUNK : Read required length of bits (Included in READ_ACCESS)
- WRITE_CHUNK: Write required length of bits (Included in WRITE_ACCESS)
- RW_CHUNK: READ_CHUNK and WRITE_CHUNK
- READ_RANDOM: READ_OFFSET and READ_CHUNK
- WRITE_RANDOM: WRITE_OFFSET and WRITE_CHUNK
- RW_RANDOM: RW_OFFSET and RW_CHUNK
- DRIVER_REUSABLE: driver is stateless and its instance can be reused safely

BZ#1170476

With this update, a completely new API that adds search capabilities for Image service and improves the performance for listing and search operations, especially on interactions with the UI is now available.

The search API allows users to execute a search query and get back search hits that match the query. The query can either be provided using a simple query string as a parameter, or using a request body. All the search APIs can be applied across multiple types within an index, and across multiple indices with support for multi index syntax.

Note: This enhancement will be removed from the Image service during the RHEL OpenStack Platform 8 (Liberty) release.

BZ#1189811

Previously, every call to policy.enforce passed an empty dictionary as the target. This prevented operators from using tenant specific restrictions in their policy.json files since the target would always be an empty dictionary. If you tried to restrict some actions so an image owner (users with the correct tenant id) could perform actions, the check categorically failed because the target is an empty dictionary.

With this update, you can pass the ImageTarget instance wrapping an

Image to the enforcer so these rules can be used and properly enforced. You can now properly grant access to the image owner(s) based on tenant (e.g., owner:%(tenant)). Without this fix, the only check that actually works in Image service is a RoleCheck (e.g., role:admin).

BZ#1198911

With this update, it is now possible to filter the list operations by more than one filter option and in multiple directions. For example:

```
/images?sort=status:asc,name:asc,created_at:desc
```

With the above, a list of images will be returned and they will be sorted by status, name, and creation date with the following directions respectively: ascending, ascending, and descending.

BZ#1201116

With this change, it is now possible to filter the list operations by more than one filter option and in multiple directions. For example:

```
/images?sort=status:asc,name:asc,created_at:desc
```

With the above, a list of images will be returned and they will be sorted by status, name, and creation date with the following directions respectively: ascending, ascending, and descending.

6.1.6. openstack-heat

BZ#1042222

The Orchestration service now includes an "OS::Heat::Stack" resource type. This OpenStack-native resource is used to explicitly create a child stack in a template. The "OS::Heat::Stack" resource type includes a 'context' property with a 'region_name' subproperty, allowing Orchestration service to manage stacks in different regions.

BZ#1053078

Resources of type AWS::EC2::SecurityGroup can now be updated in-place when their rules are modified. This is consistent with the behaviour of AWS::EC2::SecurityGroup in CloudFormation. Previously, security groups would be replaced if they were modified.

BZ#1108981

Heat now supports user hooks, which pause execution of stack operations at specified points to allow the user to insert their own actions into Heat's workflow. Hooks are attached to resources in the stack's environment file. Currently supported hook types are 'pre-create' and 'pre-update'.

BZ#1122774

The `OS::Nova::Server` resource type now includes a `'console_urls'` property. This enables the user to obtain the URL for the server's console (such as a VNC console) from the resource.

BZ#1142563

When querying a resource in the Orchestration API, a user can now request the value of one or more of the resource's attributes be included in the output. This can aid debugging, as it allows the user to retrieve data from any resource at any time without having to modify the stack's template to include that data in the outputs section.

BZ#1143805

The `OS::Cinder::Volume` resource type now includes a `'scheduler_hints'` property. This allows scheduler hints to be passed to the Block Storage service when creating a volume, and requires v2 of the Block Storage API.

BZ#1144230

The `heat-manage` command now includes a subcommand `"heat-manage service-list"`. This subcommand displays information about active `"heat-engine"` processes, where they are running, and their current status.

BZ#1149959

The `OS::Neutron::Port` resource type now supports a `'binding:vnic_type'` property. This property enables users with the appropriate permissions to specify the VNIC type of an OpenStack Networking port.

BZ#1156671

The `AWS::AutoScaling::AutoScalingGroup` resource type now supports an `'InstanceId'` property. This allows the launch configuration for an autoscaling group to be cloned from an existing server instead of an `AWS::AutoScaling::LaunchConfiguration` resource.

BZ#1159598

The `AWS::AutoScaling::LaunchConfiguration` resource type now supports an `'InstanceId'` property. This allows the launch configuration for an autoscaling group to be cloned from an existing server.

BZ#1212625

Previously, when the `'files'` section of an environment were changed in a stack update, the Orchestration service combined new files with the old stack definition to calculate the previous state. The objective of this

was to compare the previous state against the new files and new template.

As a result, the Orchestration service did not notice changes in the included files; so any updates, based solely on changes to the files, would not occur. In addition, if a previously-referenced file was removed from the environment in a stack update, the stack update would fail (though later updates with the same data could succeed).

With this release, the Orchestration service now combines the old stack with the old files to compare against the new template and new files. Updates now work as expected when editing included files in the environment.

BZ#1218692

In previous releases, changes to the absolute path of a template for a template resource (as in, a resource implicitly backed by a stack) were not recognized by the Orchestration service. This prevented nested stacks backing a template resource from being updated whenever that resource's template was renamed or moved.

With this release, the Orchestration service can now detect such changes, thereby ensuring that nested stacks are updated accordingly.

6.1.7. openstack-ironic

BZ#1151691

Bare Metal now supports the management interface of HP ProLiant Services using the iLO client python library. This allows Bare Metal to perform management operations such as retrieving/setting a boot device.

BZ#1153875

The Bare Metal service can now use cloud-init and similar early-initialization tools to insert user data on instances. Previously, doing so would have required setting up a metadata service to perform this function.

With this new update, Bare Metal can insert instance metadata onto local disk upon deployment -- specifically, to a device labeled 'config-2'. Afterwards, you can configure the early-initialization tool to find this device and extract the data from there.

BZ#1154485

The Bare Metal service can now deploy nodes using the Secure Boot feature of the UEFI (<http://www.uefi.org>). Secure Boot helps ensure that nodes boot only trusted software.

With this, the whole boot chain can be verified at boot time. You can then configure nodes to only boot authorized images, thereby enhancing security.

BZ#1154927

Bare Metal instances now feature a new field named 'maintenance_reason', which can be used to indicate why a node is in maintenance mode.

BZ#1165499

The Bare Metal service now supports Fujitsu iRMC (integrated Remote Management Controller) hardware. With this, Bare Metal can now manage the power state of such machines.

BZ#1198904

All Ironic drivers now support deployment via IPA ramdisk. IPA is written in Python, supports more features than the BASH ramdisk, and runs as a service. For these reasons, nodes deployed through IPA are generally easier to deploy, debug, and manage.

BZ#1230142

Previously, the WSMAN interface on the DRAC card would change between 11g and 12g hardware. Consequently, ``get_boot_device`` and ``set_boot_device`` calls would fail in OpenStack Bare Metal Provisioning (Ironic) when using the DRAC driver on 11g hardware. With this update, the DRAC driver checks the Lifecycle controller version, and uses alternate methods on different versions to manage the boot device. As a result, ``get_boot_device`` and ``set_boot_device`` operations succeed on 11g nodes.

BZ#1230163

The Compute service expects to be able to delete an instance at any time; however, a Bare Metal instance can only be stopped at a specific stage -- namely, when it is in the DEPLOYWAIT state. As a result, whenever the Compute service attempted to delete a Bare Metal instance that was not in the DEPLOYWAIT state, Compute's attempt failed. In doing so, the instance got stuck in a particular state, thereby required a database change to resolve.

With this release, Bare Metal instances no longer get stuck mid-deployment when Compute attempts to delete them. The Bare Metal service still won't abort an instance unless it is in the DEPLOYWAIT state.

BZ#1231327

Previously, the DRAC driver in OpenStack Bare Metal Provisioning

(Ironic) incorrectly recognized the job status 'completed with errors' as an 'in-progress' status. Consequently, ``get_boot_device`` and ``set_boot_device`` tasks failed, as they require that no in-progress jobs be present.

This update addresses this issue by adding 'completed with errors' to the list of completed statuses. As a result, ``get_boot_device`` and ``set_boot_device`` tasks will proceed even if there is a 'completed with errors' job on the DRAC card.

BZ#1231331

Previously, the ``pass_bootloader_install_info`` method was missing from the DRAC ``vendor_passthru`` interface. Consequently, PXE deployment tasks failed when local boot was enabled.

This fix adds the ``pass_bootloader_install_info`` from the standard PXE interface to ``DRAC vendor_passthru``. As a result, deployment is expected to succeed when local boot is enabled.

BZ#1233452

Prior to this update, OpenStack Bare Metal Provisioning (Ironic) operations, such as 'Power off' held a lock on a node for longer than expected.

Consequently, certain operations would fail to run while the node was still considered locked.

This update adjusts the retry timeout to two minutes. As a result, no further node lock errors have been noted.

6.1.8. openstack-keystone

BZ#1110589

The Identity Service (keystone) now allows for re-delegation of trusts. This allows a trustee with a trust token to create another trust to delegate their roles to others. In addition, a counter enumerates the number of times a trust can be re-delegated.

This feature allows a trustee to re-delegate the roles contained in its trust token to another trustee. The user creating the initial trust can control if a trust can be re-delegated when this is necessary.

Consequently, trusts can now be re-delegated if the original trust allows it.

BZ#1121844

Identity Service (keystone) now allows for unscoped tokens to be explicitly requested.

This feature was added after users who had a default project assigned were previously unable to retrieve unscoped tokens; if one of these users requested a token without defining a scope, it would be automatically scoped to the default project.

As a result of this update, unscoped tokens can now be issued to all users, even if they have a default project defined.

BZ#1165505

With this update, Identity Service (keystone), is now able to construct a hierarchy of projects by specifying a 'parent_id' within a project resource.

Previously, the Identity service only allowed for a flat project model; a project hierarchy allows for more flexible project structures, which can be used to mimic organizational structures.

As a result, Projects can now define a parent project, allowing project hierarchies to be constructed.

BZ#1189633

The Identity service now allows unscoped federation tokens to be used to obtain a scoped token using the 'token' authentication method.

When using the Identity service's federation extension, an unscoped federation token is returned as a result of the initial authentication. This is then exchanged for a scoped token. An unscoped federation token previously had to use the 'saml2' or 'mapped' authentication to obtain a scoped token. This is inconsistent with the method used to exchanging a regular unscoped token for a scoped token, which uses the 'token' method.

Exchanging an unscoped federation token for a scoped token now uses the 'token' authentication method, which is consistent with the regular unscoped token behavior.

BZ#1189639

The Identity service now restricts rescoping of tokens to only allow unscoped tokens to be exchanged for scoped tokens.

The Identity service allows for an existing token to be used to obtain a new token via the 'token' authentication method. Previously, a user with a valid token scoped for a project could use that token to obtain another token for a different project that they were authorized for. This allowed for anyone possessing a user's token to have access to any project the user has access to, as opposed to only having access to the project that the token is scoped for. To improve the security properties of scoped tokens, it was desirable to not allow this.

A new 'allow_rescope_scoped_token' configuration option is available to allow token rescoping to be restricted. Rescoping of tokens is now only allowed by using an unscoped token to authenticate when this option is enabled.

BZ#1196013

The Identity service now has an experimental support for a new token format called 'fernet'.

The token formats currently supported by the Identity service require issued tokens to be persisted in a database table. This table can grow

quite large, which requires proper tuning and a flush job to keep the Identity service performing well. The new 'fernet' token format is designed to allow the token database table to be eliminated, avoiding the problem of this table becoming a scalability limitation. The 'fernet' token format is now available as an experimental feature.

6.1.9. openstack-neutron

BZ#1108790

Prior to this update, when manually switching the tunnel source IP address on an Open vSwitch (OVS) agent, other agents kept two tunnels open to the agent: one to its old IP address and one to the new. As a result, superfluous metadata would build up on all hypervisors in the cloud running the OVS agent.

To address this, the Network node now detects a scenario where an IP address has changed on a host, persists the new information, and notifies the other agents of the IP address change.

BZ#1152579

Previously, the OpenStack Dashboard LBaaS pool details page would not correctly handle the unexpected case of the subnet attached to an LBaaS pool being deleted.

Consequently, if you created a network, subnet, router, and load balancer, and then deleted the network, subnet, and router, but retained the load balancer, the OpenStack Dashboard LBaaS details page would return error 500.

This update addresses this issue by checking for this scenario and displaying a warning message instead. As a result, the LBaaS details page now renders correctly and displays a warning as needed.

BZ#1153446

With this update, administrators are now able to view the state of High Availability routers on each node, and specifically, where the active instance is hosted.

Previously, the High Availability router state information was not previously visible to the administrator; this made maintenance harder, for example, when moving HA router instances from one agent to another, or assessing the impact of putting a node in maintenance mode.

This new functionality also serves as a sanity test and offers assurance that a router is indeed active on only one node. As a result, administrators may now run the 'neutron l3-agent-list-hosting-router <router_id>' command on a High Availability router to view where the active instance is currently hosted.

BZ#1158729

OpenStack Networking deployments with distributed routers are now able to allow tenants to create their own networks with VLAN segmentation. Previously, distributed routers only supported tunnel networks, which

may have hindered adoption as many deployments prefer to use VLAN tenant networks.

As a result of this update, distributed routers are now able to service tunnel networks as well as VLAN networks.

BZ#1213148

Red Hat Enterprise Linux OpenStack Platform 7 uses libreswan instead of openswan, however the OpenStack Networking (neutron) openswan VPNaaS driver does not function with libreswan.

With this update, you can enable the libreswan-specific driver in `vpnagent.ini`:

```
[vpnagent]
vpn_device_driver=neutron_vpnaas.services.vpn.device_drivers.libreswan_ipsec.LibreswanDrive
```

As a result, VPNaaS works as expected.

BZ#1221034

Due to a known issue with the 'python-neutron-fwaas' package, Firewall-as-a-Service (FWaaS) may fail to work. This is a result of the 'python-neutron-fwaas' package missing the database upgrade 'versions' directory.

In addition, upgrading the database schemas between version releases may not function correctly at this time.

BZ#1221076

Due to a known issue with the 'python-neutron-fwaas' package, Firewall-as-a-Service (FWaaS) may fail to work. This is a result of the 'python-neutron-fwaas' package missing the database upgrade 'versions' directory.

In addition, upgrading the database schemas between version releases may not function correctly at this time.

BZ#1227633

Previously, dnsmasq did not save lease information in persistent storage, and when it was restarted, the lease information was lost. This behavior was a result of the removal of the dnsmasq '--dhcp-script' option under BZ#1202392.

As a result, instances were stuck in the network boot process for a long period of time. In addition, NACK messages were noted in the dnsmasq log.

This update addresses this issue by removing the authoritative option, so that NAKs are not sent in response to DHCPREQUESTs to other servers. This change is expected to prevent dnsmasq from NAKing clients renewing leases issued before it was restarted/rescheduled, with the result that no DHCPNAK messages can be found in the log files.

BZ#1228096

■

In Kilo, Neutron services now can rely on so called rootwrap daemon to execute external commands like 'ip' or 'sysctl'. The daemon pre-caches rootwrap filters and drastically improves overall agent performance.

For RHEL-OSP7, rootwrap daemon is enabled by default. If you want to avoid using it and stick to another root privilege separation mechanism like 'sudo', then make sure you also disable the daemon by setting 'root_helper_daemon =' in [agent] section of your neutron.conf file.

6.1.10. openstack-neutron-lbaas

BZ#1228227

Prior to this update, the .service file was missing for the 'neutron-lbaasv2-agent' service.

Consequently, there was no way to start the agent when under control of systemd.

This update adds the missing .service file to the package.

As a result, the command 'systemctl start neutron-lbaasv2-agent' should now start the service.

6.1.11. openstack-nova

BZ#1041068

You can now use VMWare vSAN data stores. These stores allow you to use vMotion while simultaneously using hypervisor-local storage for instances.

BZ#1052804

You can now use VMware storage policy to manage how storage is assigned to different instances. This can help you ensure that instances are assigned to the most appropriate storage in an environment where multiple data stores (of varying costs and performance properties) are attached to a VMware infrastructure.

BZ#1085989

Previously, the Compute database had a missing index in the virtual_interfaces table. Because of this, as the table grew large operations on it became unacceptably long, causing timeouts.

This release adds the missing index to the virtual_interfaces table, ensuring that large amounts of data in the virtual_interfaces table do not significantly impact performance.

BZ#1193287

Support has been added for intelligent NUMA node placement for guests

that have been assigned a host PCI device. PCI I/O devices, such as Network Interface Cards (NICs), can be more closely associated with one processor than another. This is important because there are different memory performance and latency characteristics when accessing memory directly attached to one processor than when accessing memory directly attached to another processor in the same server. With this update, Openstack guest placement can be optimized by ensuring that a guest bound to a PCI device is scheduled to run on a NUMA node that is associated with the guest's pCPU and memory allocation. For example, if a guest's resource requirements fit in a single NUMA node, all guest resources will now be associated with the same NUMA node.

BZ#1203160

After fully upgrading to Red Hat Enterprise Linux OpenStack Platform 7 from version 6 (and all nodes are running version 7 code), you should start a background migration of PCI device NUMA node information from the old location to the new location. Version 7 conductor nodes will do this automatically when necessary, but the rest of the idle data needs to be migrated in the background. This is critical to complete before the version 8 release, where support for the old location will be dropped. Use 'nova-manage migrate-rhos-6-pci-device-data' to perform this transition.

Note that this is relevant only for users making use of the PCI pass-through features of Compute.

BZ#1226438

Previously, there was an error when attempting to launch an instance on a nova-network compute node configured by staypuft/openstack-foreman-installer. This was due to package contrack-tools was missing from the installer.

This bug was fixed by adding a line in openstack-nova.spec to install contrack-tools package for the nova-network's service. Nova-network can now configure networks and there is no error reported.

BZ#1228295

Previously, when the primary path to a Cinder iSCSI volume was down, a volume could not be attached to the instance, even if the Compute and Block Storage back end driver's multipath feature was enabled. This meant that users of the cloud system could fail to attach a volume (or boot a server booted from a volume).

With this fix, the host can now have a separate configuration option if the block traffic is on a separate network; the volume is then attached using the secondary path.

BZ#1229655

When deploying an OpenStack environment that uses IPv6, VNC consoles would fail to load and an exception was raised to the client because the

websocketproxy was unable to verify the origin header - "handler exception: Origin header does not match this host."

With this release, the code in websocketproxy has been updated to handle IPv6. As a result, users can now successfully connect to VNC consoles when all services are configured to use IPv6.

BZ#1230237

Previously, when attempting to evacuate a virtual machine in nova failed when used with neutron because of a failure to update port bindings. A similar issue applied to FloatingIP setup for nova-network. As a result, the virtual machine could not be evacuated because the creation of a required virtual interface failed.

With this fix, nova now correctly sets up virtual machine in both kinds of network setup. You can now evacuate virtual machines successfully.

BZ#1230485

The libvirt driver used libguestfs for certain guest inspection and modification tasks. However, libguestfs is an external library that is not updated by eventlet's monkey patch. As a result, eventlet greenthreads did not run during libguestfs API calls; this, in turn, caused the openstack-nova-compute service to hang entirely for the duration of the call. The initial call to libguestfs after installation or a system update can take seconds, during which openstack-nova-compute was unresponsive.

With this release, calls to libguestfs are now pushed to a separate, non-Eventlet threadpool. Such calls now run asynchronously, and do not impact the responsiveness of openstack-nova-compute.

BZ#1242502

Previous releases used incorrect data versioning, which caused the PCI device data model to be sent in an incorrect format. This, in turn, prevented the openstack-nova-compute service from starting if there were any PCI-passthrough devices whitelisted.

This release now uses correct data versioning, thereby allowing openstack-nova-compute to start and register any whitelisted PCI devices.

6.1.12. openstack-packstack

BZ#1185652

This feature adds IPv6 support to Packstack, allowing Packstack to use IPv6 address as values in networking-related parameters such as CONFIG_CONTROLLER_HOST, CONFIG_COMPUTE_HOSTS, and CONFIG_NETWORK_HOSTS.

6.1.13. openstack-puppet-modules

BZ#1231918

Previously, puppet-neutron did not allow for customization of the neutron dhcp_domain setting. As a consequence, the overcloud nodes would be offered an invalid domain suffix by the undercloud DHCP. With this update, the neutron dhcp_domain setting has been made configurable, and defaults to an empty domain suffix.

BZ#1236057

Previously, the HAProxy configuration of the Telemetry service used incorrect checks, which caused the Telemetry service to fail in an HA deployment. Specifically, the HAProxy configuration did not have availability checks, and incorrectly used SSL checks instead of TCP.

This release fixes the checks, ensuring that the Telemetry service is correctly balanced and can launch in an HA deployment.

BZ#1244358

The Director uses misconfigured HAProxy settings when deploying the Bare Metal and Telemetry services with SSL enabled in the undercloud. This prevents some nodes from registering.

To work around this, comment out 'option ssl-hello-chk' under the Bare Metal and Telemetry sections in /etc/haproxy/haproxy.cfg after installing the undercloud.

6.1.14. openstack-sahara

BZ#1149055

This enhancement adds namenode high availability as a supported option in the HDP 2.0.6 plugin.

Users can signal that they require a cluster to be generated in HA mode, by passing a cluster with a quorum of zookeeper servers and journalnodes, and at least 2 namenodes. For example:

```
"cluster_configs": {
  "HDFSHA": {
    "hdfs.nnha": true
  }
}
```

BZ#1155378

With this enhancement, the Sahara API now fully supports the HTTPS protocol.

BZ#1158163

-

Prior to this update, Sahara's 'distributed' mode feature was in alpha testing. Consequently, Red Hat Enterprise Linux OpenStack Platform did not package or support the 'sahara-api' or 'sahara-engine' processes individually.

With this update, the 'distributed' mode feature is considered stable, and RHEL OpenStack Platform now provides systemd unit files for the 'sahara-api' and 'sahara-engine' services.

As a result, users can run Sahara in distributed mode, with separation of the API and engine node clusters.

BZ#1164087

Sahara objects can now be queried by any field name. This is done using the GET parameters that match the API field names, as seen on list methods.

BZ#1189500

This enhancement adds a CLI that allows configuration of the default cluster templates for each major plugin. The provision of default templates is expected to speed and facilitate end-user adoption of Sahara.

As a result of this update, administrators can now add shared default templates for adaptation and direct usage by customers.

BZ#1189504

Integration tests for Sahara have been refactored from more brittle pure python tests to allow easy, YAML-based configuration to define "scenarios".

BZ#1189511

Previously, the cm_api library was not packaged by Cloudera for any Linux distribution. The previous CDH plug-in depended on this package, so CDH could not be enabled as a default plug-in prior to this release. Now, a subset of the cm_api library has been added to Sahara's codebase, and CDH is functional and enabled by default.

BZ#1192290

Previously, many of the processes in cluster creation polled infinitely. Now, timeouts have been added for many stages of cluster creation and manipulation, and users are shown appropriate error messages when cluster operations have taken longer than is reasonable.

BZ#1194532

A new endpoint has been added to Sahara that allows queries of the available job types per plug-in and version that the Sahara installation supports. This information is useful both for UI presentation and

filtering, and for CLI and REST API users.

BZ#1214817

Prior to this release, Red Hat Enterprise Linux OpenStack Platform did not package or support the sahara-api or sahara-engine processes individually, because Sahara's "distributed" mode was in alpha testing. Now that this feature is stable, RHEL OpenStack Platform provides systemd unit files for the sahara-api and sahara-engine services, and users can use Sahara in distributed mode, with separation of api and engine node clusters.

BZ#1231923

Previously, the HDP plug-in installed the Extra Packages for Enterprise Linux (EPEL) repository on cluster generation, even though neither the plug-in nor the saraha-image-elements package used the repository for any purpose. Consequently, a needless, potentially error-prone step was introduced into HDP cluster generation, and on update these clusters might update with unsupported packages. Now, the repository is no longer installed by the HDP plug-in.

BZ#1231974

A logrotate file that enforces size limitations within the current Red Hat OpenStack standard has been added to prevent log files from becoming too large before they are rotated.

BZ#1238700

Prior to this update, while NameNode HA for HDP was functional and feature complete upstream, Sahara continued to point Oozie at a single NameNode IP for all jobs. Consequently, Oozie and Sahara's EDP were only successful when a single, arbitrary node was designated active (in an A/P HA model). This update addresses this issue by directing Oozie to the nameservice, rather than any one namenode. As a result, Oozie and EDP jobs can succeed regardless of which NameNode is active.

6.1.15. openstack-selinux

BZ#1233154

Prior to this update, Neutron was trying to bind to port that it was not allowed to use. Consequently, SELinux prevented Neutron from working. Now, Neutron is allowed to connect to unreserved ports and runs without issues.

BZ#1240647

■

Previously, the Neutron VPN agent was started with the wrong context. As a consequence, SELinux prevented the VPN agent from running. With this update, the Neutron VPN agent has the proper context, and as a result, it is able to run in enforcing mode.

6.1.16. python-django-horizon

BZ#1101375

OpenStack Trove instances can now be resized in the OpenStack dashboard user interface by selecting a new flavor for the instance.

BZ#1107490

The 'API Access' page in the dashboard ('Project > Compute > Access & Security > API Access') now provides more information on user credentials. To view this information, click 'View Credentials'. A pop-up displays the user name, project name, project ID, authentication URL, S3 URL, EC2 URL, EC2 access, and secret key.

BZ#1107924

The option to create Block Storage (cinder) volume transfers has been added to the 'Volumes' tab in the OpenStack dashboard. Volume transfers move ownership from one project to another. A donor creates a volume transfer, captures the resulting transfer ID and secret authentication key, and passes that information out of band to the recipient (such as by email or text message). The recipient accepts the transfer, supplying the transfer ID and authentication key. The ownership of the volume is then transferred from the donor to the recipient, and the volume is no longer visible to the donor.

Note the following limitations of the Block Storage API for volume transfers and their impact on the UI design:

1. When creating a volume transfer, you cannot specify who the intended recipient will be, and anyone with the transfer ID and authentication key can claim the volume. Therefore, the dashboard UI does not prompt for a recipient.
2. Current volume transfers are only visible to the donor; users in other projects are unable to view these transfers. So, the UI does not include a project table to view and accept volume transfers, since the current transfers are not visible. Instead, the transfer information is added to the volume details, which are visible by the donor, and the volume state clearly reflects that a transfer has been created. The UI also cannot present to the recipient a pull-down list of transfers to accept.
3. The only time that the authorization key is visible to the donor is in the response from the creation of the transfer; after creation, it is impossible for even the donor to recover it. Since the donor must capture the transfer ID and authorization key in order to send it to the recipient, an extra form was created to present this information to the donor immediately after the transfer has been created.

BZ#1112481

OpenStack Dashboard now uses Block Storage (cinder) version 2 as its preferred version. Now when a Block Storage client is requested, access is given using cinder version 2, if not specified otherwise.

BZ#1114804

You can now use the dashboard to view, import, and associate metadata definitions that can be used with various resource types (images, artifacts, volumes, flavors, aggregates, etc).

BZ#1121848

In OpenStack Dashboard, the instance detail page now displays the host node. This data is intended to assist when diagnosing issues.

BZ#1124672

This update adds partial support for Domain Admins to the OpenStack Dashboard. In addition, when using Identity Service (keystone) version 3, a newly-created user does not need to have a primary project specified.

BZ#1143807

You can now disable and enable compute hosts through the dashboard. This capability is available through the 'Actions' column of every compute host in 'Admin > Hypervisors > Compute Host'.

Disabling a compute host prevents the scheduler from launching instances using that host.

BZ#1150839

The 'Manage/Unmanage' option has been added to the 'Volumes' tab of the OpenStack dashboard. 'Manage' takes an existing volume created outside of OpenStack and makes it available. 'Unmanage' removes the visibility of a volume within OpenStack, but does not delete the actual volume.

BZ#1156678

The user interface options available in the dashboard for the OpenStack Orchestration service (heat) have been improved. For example, users can now check, suspend, resume, and preview stacks.

BZ#1162436

The results displayed in tables for the Data Processing service can now be filtered to allow the user to see only those results that are

relevant.

BZ#1162961

You can now flag a volume as 'Bootable' through the dashboard.

BZ#1166490

The OpenStack dashboard can now use a custom theme. A new setting, 'CUSTOM_THEME_PATH' was added to /etc/openstack_dashboard/local_settings file. The theme folder should contain one _variables.scss file and one _styles.scss file. The _variables.scss file contains all the bootstrap and Horizon-specific variables that are used to style the graphical user interface, and the _styles.scss file contains extra styling.

BZ#1170470

SRIOV can now be configured in the OpenStack dashboard. Options include exposing further information on the 'Port Details' tab, and allowing port type selection during port creation and update.

BZ#1170471

This enhancement allows you to view encryption metadata for encrypted volumes in OpenStack Dashboard (horizon). A function to display encryption metadata was added, and allows the user to click on the "Yes" in the Encrypted column, and be taken to a page where the encryption metadata is visible.

BZ#1186380

When uploading an image through the dashboard, you can now select OVA as its format. In previous releases, OVA was not available as an option.

BZ#1189711

The dashboard now provides wizards for creating and configuring the necessary components of the OpenStack Data Processing feature. These wizards are useful for guiding users through the process of cluster creation and job execution. To use these wizards, go to 'Project > Data Processing > Guides'.

BZ#1189716

This enhancement adds ceilometer IPMI meters to OpenStack Dashboard. Six ipmi meters have been exported from ceilometer; the methods 'list_ipmi' and '_get_ipmi_meters_info' are used to retrieve the meter data.

BZ#1190312

You can now view details about Orchestration service hosts through the dashboard. To do so, go to 'Admin > System > System Information > Orchestration Services'. This page is only available if the Orchestration service is deployed.

6.1.17. python-glance-store

BZ#1236055

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.

6.1.18. python-ironicclient

BZ#1212134

Previously, certain operations in OpenStack Bare Metal Provisioning (Ironic) would fail to run while the node was in a `locked` state. This update implements a `retry` function in the Ironic client. As a result, certain operations take longer to run, but do not fail due to `node locked` errors.

6.1.19. python-openstackclient

BZ#1194779

The python-openstackclient package is now re-based to upstream version 1.0.3. This re-base features new fixes and enhancements relating to support for the Identity service's v3 API.

6.1.20. qemu-kvm-rhev

BZ#1216130

On a virtual disk with a high number of sectors, the number of sectors was in some cases handled incorrectly, and converting a QEMU image failed with an "invalid argument" error. This update fixes the incorrect calculation that caused this error, and the described failure no longer occurs.

BZ#1240402

Due to an incorrect implementation of portable memory barriers, the QEMU emulator in some cases terminated unexpectedly when a virtual disk was under heavy I/O load. This update fixes the implementation in order to

achieve correct synchronization between QEMU's threads. As a result, the described crash no longer occurs.

6.1.21. sahara-image-elements

BZ#1155241

This package allows users to create HDP 2.0.6 and CDH 5.3.0 images for use in RHEL OpenStack Platform 7.

BZ#1231934

Previously, CDH image generation sometimes failed, because the image creation wrapper script specified too small a space for generation of the CDH image on some systems. Now, the image generation space is increased for CDH images, and images are generated successfully.

6.1.22. sos

BZ#1232720

When using the sosreport utility on a Pacemaker node, one of the MariaDB MySQL server log-files was not properly collected. With this update, the underlying code has been corrected, and the log-file is now collected as expected.

BZ#1240667

Previously, various OpenStack plug-ins for the sosreport utility were incorrectly collecting passwords in plain text. As a consequence, the compressed file created after using sosreport could contain human-readable passwords. This update adds obfuscation of all passwords to sosreport OpenStack plug-ins, and the affected passwords in the sosreport tarball are no longer human-readable.

6.2. RHEA-2015:1549 — RED HAT ENTERPRISE LINUX OPENSTACK PLATFORM DIRECTOR RELEASE

The bugs contained in this section are addressed by advisory RHEA-2015:1549. Further information about this advisory is available at <https://access.redhat.com/errata/RHEA-2015:1549.html>.

6.2.1. diskimage-builder

BZ#1230823

A missing dependency on PyYAML from diskimage-builder caused image builds to fail. This fix adds the dependency. Now image builds are successful.

6.2.2. instack

BZ#1210479

The `deploy-baremetal-overcloudrc` file is no longer used in the deployment process. Deployments now use the "openstack overcloud deploy" command and the associated command line arguments.

6.2.3. instack-undercloud

BZ#1205825

A duplicate line in the RabbitMQ configuration caused the Undercloud installation to fail. This fix removes the duplicate line and the Undercloud installation is now successful.

BZ#1229296

The `openstack-ironic-discoverd` service would check for `openstack-ironic-api`'s presence on start up. Due to start up ordering, `openstack-ironic-discoverd` failed to detect `openstack-ironic-api` on reboot. This fix removes the need to check for `openstack-ironic-api`'s presence from `openstack-ironic-discoverd` start up code. All services start successfully now on reboot.

6.2.4. openstack-ironic

BZ#1190481

A transient error in the Ironic/Heat `systemd` service file incorrectly redirected the logs. This fix corrects the Ironic/Heat service file and the logs redirect to the correct file.

6.2.5. openstack-ironic-discoverd

BZ#1227755

The `edeploy` plugin for `ironic-discoverd` collected too much information to store in a SQL blob. Discovery failed when `edeploy` data was posted to Ironic because the column would overflow. This fix changes the `edeploy` plugin to store data in a Swift object on the Undercloud. Discovery no longer fails when using the `edeploy` plugin.

6.2.6. openstack-tripleo-heat-templates

BZ#1232269

Hostnames were inappropriately configured on the Overcloud nodes. This

meant Pacemaker could not resolve the cluster member's name. This fix shortens the hostnames, which means Nova and Neutron now do not append invalid domain names.

BZ#1232439

With this release, a new enhancement allows allocating specific IP address ranges for isolated networks.

As a result, it is now possible to specify IP address allocation by setting the following parameters in the Overcloud Orchestration template:

- * InternalApiAllocationPools
- * StorageAllocationPools
- * StorageMgmtAllocationPools
- * TenantAllocationPools
- * ExternalAllocationPools

BZ#1232461

With this release, you can now configure the Compute VNC proxy network. This allows you to isolate various networks and specify which services run where.

As a result, it is now possible to specify which network to use for the Compute VNC proxy by setting the ServiceNetMap -> NovaVncProxyNetwork parameter in the Overcloud Orchestration template.

BZ#1232485

This enhancement add VLAN identifiers and OVS bond options to Heat templates. This reduces the amount of duplicate configuration. The Heat templates now contain the following parameters: BondInterfaceOvsOptions, StorageNetworkVlanID, StorageMgmtNetworkVlanID, InternalApiNetworkVlanID, TenantNetworkVlanID, ExternalNetworkVlanID,

BZ#1232747

HAProxy failed to configure the Horizon listener, which caused unavailability of Horizon on the public VIP. This fix enables the HAProxy Horizon listener and the Horizon dashboard is now available on the public VIP.

BZ#1232797

Ceilometer uses an incorrect redis VIP for its backend_url. This fix sets the Ceilometer backend_url to one that Heat provides and not constructed during deployment. Ceilometer now uses the correct IP address.

BZ#1232938

The novncproxy service failed to start due to socket binding conflicts. This meant novncproxy service would not be available. This fix configures the novncproxy service to bind only on the controller internal_api address. The novncproxy now starts successfully.

BZ#1233061

Previously, a race condition occurred during the initialization of the neutron database when neutron-server was first run. This error was seen when two controllers happened to start neutron-server simultaneously. Subsequently, the startup of neutron-server and agents failed on the controller node that lost the race, and as a consequence, Neutron services failed to start on the affected controller nodes. Errors in the logs look like the following:

```
DBDuplicateEntry: (IntegrityError) (1062, "Duplicate entry
'datacentre-1' for key 'PRIMARY'") 'INSERT INTO ml2_vlan_allocations
(physical_network, vlan_id, allocated) VALUES (%s, %s, %s)'
(('datacentre', 1, 0),
```

With this release, the Neutron server is momentarily started and then stopped on one node, the pacemaker master, allowing this initial database setup to happen, before allowing the rest of the puppet or pacemaker configuration to happen. As a result, Neutron services are brought up on all controllers nodes without error.

BZ#1233283

The mongodb node list would only build correctly on one Controller node, which meant all other Controller nodes had Ceilometer configured with an empty list of mongodb nodes. This fix corrects the mongodb node list on all Controller nodes. Ceilometer has a properly populated list of mongodb nodes.

BZ#1234637

An incorrect configuration of the HAProxy backend caused HAProxy to forward requests for the glance-registry service to offline nodes. This fix monitors the glance-registry service for updates to ensure offline nodes are detected.

BZ#1234817

The HAProxy listener for Galera would bind on the ctlplane address. This meant clients could not reach the Galera service when using an Overcloud with network isolation. This fix changes the binding address of the HAProxy Galera listener to the VIP in the internal_api network. Clients now can reach the Galera service on Overclouds with network isolation.

BZ#1235408

HAProxy did not use clustercheck to check MariaDB's backends status.

This caused HAProxy to forward requests to MariaDB nodes responsive at the TCP check but not in synchronization with the Galera cluster. This fix now uses clustercheck to check MariaDB's backends status. HAProxy now forwards requests to MariaDB nodes correctly.

BZ#1235421

The distro-specific hieradata file did not apply onto the Overcloud nodes. This fix provides the static RedHat.yaml for distribution on all Overcloud nodes.

BZ#1235454

The mariadb service started on boot, which caused Pacemaker's mariadb resource to fail after a reboot. This fix disables the mariadb service from automatically starting on boot. This means mariadb is fully controlled as a Pacemaker resource.

BZ#1235703

The Keystone service started through Pacemaker and as a systemd dependency of the Ceilometer resource in Pacemaker. This caused conflicts between the two versions of the Keystone service starting that produced failures for the Pacemaker Keystone resource to start. This fix adds a Pacemaker constraint to halt the Ceilometer resource until the Keystone resource starts. Keystone starts before Ceilometer and the Keystone service does not start through systemd.

BZ#1235848

Deployments with the director plans behaved differently to Heat template-based deployments due to a missing ControlPlaneNetwork parameter. This caused plan-based Overcloud deployments to fail when using network isolation. This fix includes a patch to include the ControlPlaneNetwork parameter. Now the Overcloud deploys properly.

BZ#1236374

Heat services restarted on unrelated redis VIP relocation. In Pacemaker, the Heat resource failed to restart due to dependencies on the Ceilometer resource, which failed to restart on relocation of the redis VIP due to clustering failures. This fix stops Heat from restarting when Ceilometer restarts.

BZ#1236407

On Overclouds with network isolation enabled, Pacemaker set the redis master to a hostname on a network where the master was unreachable. This meant redis nodes failed to join the cluster. This fix resolves Pacemaker hostnames against the internal_api addresses when deploying with network isolation.

BZ#1238117

Previously, OpenStack was using the NeutronScale puppet resource that was enabled on controller nodes and tasked with rewriting the neutron agents' "host" entries to look like "neutron-n-0" on controller 0 or "neutron-n-1" on controller 1. This renaming was done toward the end of the deployment, when the corresponding neutron-scale resource was started by pacemaker. Mostly reported in VM environments, neutron would subsequently complain about not having enough L3 agents for L3 HA, and there would be inconsistency in the overcloud neutron agent-list. Consequently, in some cases, the error manifested itself in an error message from Neutron that there were not enough L3 agents to provide HA (the default minimum of 2). The "neutron agent-list" command on the overcloud would show inconsistency in the agents; for example, duplicate entries for each agent with both the original agent on host "overcloud-controller-1.localdomain" (typically shown "XXX") and the "newer" agent on host "neutron-n-1" (alive status ":-"). In other cases, agent renaming would cause one of the neutron agents, openvswitch, to fail when there was only one controller, and then the rest of the agents under it would also fail to start as they were chained, resulting in no L3, metadata, or dhcp agents.

This problem has been fixed by ensuring that the native neutron L3 High Availability is used, and that enough DHCP agents per network are enabled for native neutron HA. The latter is a needed addition as it was previously statically set at two in all cases. This was added as a configurable parameter in the tripleo heat templates with a default value of '3' and also wired up to deploy in the oscplugin. The NeutronScale resource itself has been removed from the tripleo heat templates where the overcloud controller puppet manifest is kept. As a result, deployments made after this fix will not have the neutron-scale resource on controller nodes, which can be verified by the following commands:

```
1. On a controller node:
# pcs status | grep -n neutron -A 1
```

You should not see any "neutron-scale" clone set or resource definition.

```
2. On the undercloud:
$ source overcloudrc
$ neutron agent-list
```

All the neutron agents should be reported as being on a host with a name like "overcloud-controller-0.localdomain" or "overcloud-controller-2.localdomain" but not "neutron-n-0" or "neutron-n-2".

BZ#1238336

Controller nodes did not share consoleauth tokens, which caused failures with parts of authentication requests. This fix incorporates memcached to share consoleauth tokens. Authentication requests are now successful.

BZ#1240631

■

This release allows you to configure the `neutron_tunnel_id_ranges` and `neutron_vni_ranges` parameters, which govern the GRE or VXLAN tunnel IDs respectively. These are to be made available for overcloud tenant networks by overcloud Neutron. As an example, you can specify:

```
# openstack overcloud deploy --plan overcloud --control-scale 3 --
compute-scale 1 --neutron-tunnel-id-ranges "1:1111,2:2222" --neutron-
vni-ranges "3:33,5:55,100:999" --neutron-tunnel-types "gre,vxlan"
```

If not specified, `tunnel_id_ranges` and `neutron_vni_ranges` both default to "1:1000", which may be unsuitable or restrictive for some deployment scenarios.

If deploying as shown in the example above, you can inspect and verify the relevant neutron configuration files on a controller node (post deploy), for instance:

```
# grep -rni 'vni_ranges\|id_ranges\|tunnel_types' /etc/neutron/*
/etc/neutron/plugin.ini:78:tunnel_id_ranges =1:1111,2:2222
/etc/neutron/plugin.ini:85:vni_ranges =3:33,5:55,100:999
/etc/neutron/plugins/openvswitch/ovs_neutron_plugin.ini:77:tunnel_types
=gre,vxlan
/etc/neutron/plugins/ml2/ml2_conf.ini:78:tunnel_id_ranges =1:1111,2:2222
/etc/neutron/plugins/ml2/ml2_conf.ini:85:vni_ranges =3:33,5:55,100:999
```

You can also test by creating an overcloud network (again, post-deploy), for instance:

```
# source overcloudrc # from the undercloud box for example
# neutron net-create --provider:network_type "vxlan" "foo"
```

This will create a vxlan type network. On inspection you can verify that this received a segmentation ID from the specified `vni_ranges`:

```
# neutron net-show foo
| provider:network_type      | vxlan
| provider:segmentation_id  | 3
```

Likewise, you can verify an appropriate segmentation ID is assigned for GRE networks.

BZ#1241231

Previously, uncoordinated creation of the overcloud Keystone admin tenant caused errors in a deployment with more than one controller. As a consequence, Heat stack creation failed on the `ControllerNodesPostDeployment` resource, and Keystone returned a 409 ERROR: `openstack Conflict occurred attempting to store project - Duplicate Entry (HTTP 409)`. The Puppet run failed at that point. The fix for this problem includes the creation of the admin tenant on the pacemaker master node first. As a result, deployments with more than one controller correctly create the overcloud keystone admin user. After a successful deployment you can verify this by interacting with any of the overcloud services as the admin user:

```
# on the undercloud system, for example:
```

```
$ source overcloudrc
$ keystone user-list
```

BZ#1241610

Tuskar modified the names of the top level parameters in the Heat stack when deploying an Overcloud. This caused an error during stack validation from Heat:

```
ERROR: The Parameter (NeutronExternalNetworkBridge) was not defined in
template.
```

As a workaround, use "tuskar plan-update" to modify the parameter, or use the modified parameter name in the environment file:

```
parameters:
  Controller-1::NeutronExternalNetworkBridge: ""
```

Overcloud deploys using the correct parameter value.

Note: the parameter needs to be defined in the "parameters:" section, not the "parameter_defaults:" section. Otherwise, the value set in Tuskar's exported environment.yaml overrides the value.

BZ#1244013

Compute nodes queried Keystone for the Cinder publicurl endpoint, regardless of whether they had connectivity. This meant dedicated Compute nodes failed to interact with Cinder API. This fix changes the publicurl endpoint to the internalurl endpoint, which Compute nodes can access.

BZ#1244019

Cinder Storage nodes queried Keystone for the Nova and Swift publicurl endpoints, regardless of whether they had access. This meant dedicated Cinder Storage nodes could not interact with Nova and Swift APIs. This fix changes the publicurl endpoints to internalurl endpoints, and Cinder Storage nodes now communicate with Nova and Swift APIs successfully.

BZ#1244226

Prior to this release, the allow_overlapping_ips Neutron setting was left at default. As a consequence, allow_overlapping_ips in Neutron was disabled, preventing definition of multiple networks with the same ranges. This problem has been fixed by setting allow_overlapping_ips to "true", and as a result, networks with overlapping ranges can be defined.

6.2.7. openstack-tripleo-image-elements

BZ#1235994

Previously, the default Ceph scale was set to 1, resulting in a Ceph node being created even if the user did not want to create one.

With this release, the default Ceph scale setting has been changed to 0. As a result, Ceph is not deployed by default anymore.

6.2.8. openstack-tripleo-puppet-elements**BZ#1229302**

The `os-apply-config` command created `/etc/puppet/hieradata` with open permissions. The files in this directory contained passwords and tokens that could provide unauthorized access to the OpenStack installation. This fix sets `/etc/puppet/hieradata` as a root-owned directory with `0700` permissions. Only the root user can access `/etc/puppet/hieradata`, which provides a more secure installation.

BZ#1233916

Overcloud nodes contained incorrectly synchronized system times. This resulted in various errors across the HA Controller cluster. As a workaround, pass the `--ntp-server` command line argument when running the "openstack overcloud deploy" command. This argument configures the ntp server in `/etc/ntp.conf` on each Overcloud node and starts the ntpd service. This produces properly synchronized system times and a successful Overcloud deployment.

6.2.9. openstack-tuskar**BZ#1205281**

Previously, migration to a puppet deployment meant that the boot-stack tripleo-image-element was no longer creating the initial Tuskar database. Consequently, after successful installation of the undercloud, the Tuskar service was not correctly configured. This release ensures that Tuskar is properly installed and configured for the undercloud. As a result, after successful installation of the undercloud, you can interact with the Tuskar service.

BZ#1220651

The tuskar service configuration parameter `auth_strategy` defaulted to "noauth". This allowed unrestricted access to the tuskar management plan and roles, including templates and any set sensitive parameters like passwords. This fix sets the default to keystone authentication. Now non-authenticated http requests to the tuskar service will return a HTTP 401 Unauthorized error. Use the following command to verify from the

Undercloud:

```
$ curl -v localhost:8585/v2/plans
```

6.2.10. openstack-tuskar-ui

BZ#1197857

The code for executing bulk actions did not include a check to ensure a node is actually selected. The actions assumed that the list of nodes was not empty. The code generated an uncaught exception, resulting in a "Something went wrong" message when a DEBUG is disabled.

With this update, a check to verify if at least one node is selected is included or to display a helpful error message when no node is selected. As a result, trying to perform a bulk action with no nodes selected results in a helpful message telling you to select some nodes first.

BZ#1227013

The node detail URL name was specified incorrectly and rendered incorrectly. This fix corrects the node detail URL name. Now the node detail link renders correctly, which allows access to the node detail page.

BZ#1232329

Roles were not added to the plan during Undercloud installation. As a result, there could not be edited using the OpenStack Dashboard.

With this update, the Undercloud installation has been fixed to add roles to the plan. As a result, you can now edit these roles using the Dashboard.

BZ#1236360

The director's user interface took long periods with page loads due to communicating with External API services such as keystone, heat, ironic and tuskar. This fix adds caching to all External service calls to reduce the amount of calls and decrease page load times. Page loading times are now significantly faster when accessing the user interface.

BZ#1245192

The user interface attempted to connect to the Overcloud keystone-api before the creation and initialization of the corresponding endpoint. The user interface would not find the endpoint the Overcloud would not initialize from the user interface. This fix enables the user interface to properly identify when the Overcloud has not yet been initialized. The error no longer occurs.

6.2.11. python-rdomanager-oscplugin

BZ#1229795

The post-install validation was not implemented in the unified CLI, it was implemented only in the deprecated OpenStack Deployment (tripleO) scripts. As a result, the command was missing from python-rdomanager-oscplugin.

With this update, the post-install validation is implemented and the 'openstack overcloud validate' command is now present in the CLI.

BZ#1229796

The DRAC ready-state configuration was not implemented in the unified CLI, it was implemented only in the deprecated OpenStack Deployment (tripleO) scripts. As a result, the command was missing from python-rdomanager-oscplugin.

With this update, the DRAC ready-state is implemented and 'openstack baremetal configure ready state' command is now present in the CLI.

BZ#1230450

This release introduces four parameters that govern the default overcloud neutron tenant network behavior: `--neutron-network-type`, `--neutron-tunnel-types`, `--neutron-disable-tunneling`, and `--neutron-network-vlan-ranges`. The default is to provide GRE types and tunnels for overcloud tenant networks.

If the `--neutron-network-type` is specified as 'vlan', you can also set the `--neutron-network-vlan-ranges` parameter, which governs the range of VLAN IDs to be allocated to overcloud Neutron tenant networks (defaults to "datacenter:1:1000", where "datacenter" must be the name of the Neutron physical network where the VLANs will be allocated). If exclusively specifying 'vlan' as the network type, then `--neutron-disable-tunneling` must also be specified (examples below).

If the `--neutron-network-type` is specified as 'gre' or 'vxlan' (or both), then the corresponding `--neutron-tunnel-types` must also be enabled (examples below). Furthermore, when specifying the 'gre' or 'vxlan' network types, you can further set the related `--neutron-tunnel-id-ranges` and `--neutron-vni-ranges` parameters that govern the GRE or VXLAN tunnel IDs respectively. These are to be made available for the tenant networks.

Examples:

* Specify the 'vxlan' type for overcloud tenant networks:

```
openstack overcloud deploy --plan overcloud --control-scale 3 --compute-scale 2 --neutron-network-type "vxlan" --neutron-tunnel-types "vxlan"
```

* Specify both 'gre' and 'vxlan' overcloud tenant networks:

```
openstack overcloud deploy --plan overcloud --control-scale 3 --compute-scale 2 --neutron-network-type "gre,vxlan" --neutron-tunnel-types "gre,vxlan"
```


* Specify the 'vlan' type for overcloud tenant networks, as well as the VLAN ranges:
 openstack overcloud deploy --plan overcloud --control-scale 3 --compute-scale 2 --neutron-network-type "vlan" --neutron-disable-tunneling --neutron-network-vlan-ranges "datacenter:3:3000"

These parameters default to the following values when not specified at deployment time, and may be deemed restrictive or unsuitable for a given deployment scenario:

```
--neutron-network-type # default 'gre'
--neutron-tunnel-types # default 'gre'
--neutron-disable-tunneling # defaults to tunneling being enabled
--neutron-network-vlan-ranges # 'datacenter:1:1000'
```

After deployment, you can verify that the following are set in the Neutron configuration logs on a controller node; the values here should correspond to those specified during the deployment as shown above:

```
grep -rni 'tunnel_types\|network_type\|enable_tunneling\|vlan_ranges'
/etc/neutron/*
/etc/neutron/plugin.ini:14:tenant_network_types = vlan
/etc/neutron/plugin.ini:72:network_vlan_ranges =datacenter:3:3000
/etc/neutron/plugins/openvswitch/ovs_neutron_plugin.ini:54:enable_tunneling=False
/etc/neutron/plugins/ml2/ml2_conf.ini:14:tenant_network_types = vlan
/etc/neutron/plugins/ml2/ml2_conf.ini:72:network_vlan_ranges
=datacenter:3:3000
```

Also, you can create a network in overcloud Neutron and inspect it; for example, on the undercloud:

```
$ source overcloudrc
$ neutron net-create "default"
$ neutron net-show default
...
| provider:network_type      | vlan          |
| provider:physical_network | datacenter    |
| provider:segmentation_id  | 3             |
...
```

BZ#1232851

The Unified CLI displays a timed out error, expecting the OpenStack overcloud deployment to have been incomplete. As a result, the user incorrectly gets a message that the resulting deployment failed.

With this update, the timeout value is increased to an hour. As a result, deployments should now finish within the timeout.

BZ#1233201

With this release, a duplicated command for scaling a deployment has been removed in favor of scaling a deployment directly using Tuskar or using the deploy command.

BZ#1234343

Issues in the KVM PXE code displayed failures when too many nodes tried to PXE-boot simultaneously, resulting in some nodes failing to connect to DHCP.

With this update, the sleep value is increased, allowing introspection on the nodes. As a result, DHCP is no longer an issue, making the introspection a little longer.

BZ#1234673

Previously, when using the 'openstack overcloud deploy' command, some of the command line parameters were not being used by the code. As a result, the user was unable to set these parameters.

With this release, the parameters are now used and sent to the Orchestration service.

BZ#1236073

The OpenStack Overcloud validate ran the OpenStack Integration Test Suite (tempest) as documented upstream instead of using a midstream tool in redhat-openstack/tempest, resulting in an output format that was different.

With this update, the OpenStack Overcloud validate runs tempest using the midstream tool and the output format is as expected.

BZ#1236399

The CLI had no method of customizing certain Neutron parameters. This fix adds new arguments to the deploy command. The user provides these arguments with "Openstack overcloud deploy" command to configure these parameters.

BZ#1237068

Tuskar API expected all parameters to be strings, but since some of the parameters were passed as integer values, the API would reject those.

With this update, all parameters are converted to strings before being sent to the API. As a result, the API accepts all parameters and the deployment works successfully.

BZ#1240101

The director used hard coded paths to the Heat template collection in `/usr/share/openstack-tripleo-heat-templates/`. This caused problems when using a custom Heat template e.g. creating a local copy and customizing it. This fix provides variable input with the `--template` argument. Users can now create an Overcloud with a customised Heat template collection.

BZ#1240461

Previously, the necessary parameters were not passed to Orchestration service correctly for an OpenStack Overcloud redeployment, resulting in an `UPGRADE_FAILED` message.

With this update, a parameter `'Existing=True'` is added to the Orchestration update parameters, resulting in a successful Overcloud deployment.

BZ#1240679

During deployment, the heat engine logs return a non-zero error due to no valid hosts being found, despite the ironic logs showing nodes available. This is due to the director not setting nodes to "available" when the "openstack baremetal introspection" command completes. This fix sets the nodes to "available" after the introspection completes. The director now sees the nodes when deploying the Overcloud.

BZ#1242967

In addition to using plans stored in its database, the director now also uses Heat template collections to create an Overcloud. This provides a method to use Heat templates directly when running "openstack overcloud stack update" command to update packages on the Overcloud. If you created the Overcloud using a template collection instead of a plan, run "openstack overcloud stack update" command with the `--templates` parameter instead of the `--plan` parameter.

BZ#1242973

The director now provides parameters to accept extra Heat environment files when updating packages on Overcloud nodes. Overcloud isolated networks require extra environment files upon creation and subsequent updates. Without the extra environment files, users cannot use "openstack overcloud update stack" on an Overcloud using isolated network. User now use the `-e` parameter with the "openstack overcloud update stack" command to update packages on Overcloud nodes using isolated networks.

BZ#1242989

This release allow the use of Heat templates directly when deleting overcloud nodes. You can now run the "openstack overcloud node delete" command with the `--templates` parameter if the overcloud was deployed without Tuskar.

BZ#1242990

New in this release is the ability to accept extra Heat environment files when deleting a node from overcloud. When using an isolated network, extra environment files must be passed to Heat on any overcloud update. With this update, the "openstack overcloud node delete" command can be used to delete overcloud nodes on an isolated network.

BZ#1243016

Some arguments existed that provided the same function as other arguments for the deploy command. This fix removes the redundant arguments (`--plan-uuid`, `--use-tripleo-heat-templates`, `--extra-template`).

BZ#1243365

A timeout could not be set, which meant deployments always timeout after one hour. This fix adds a timeout argument that allows users to set a custom timeout. It defaults to four hours.

BZ#1243846

The Heat templates lacked a certain parameter to configure the Neutron L3 Agent. As a result the director would not configure the L3 Agent properly. This fix adds the missing Heat parameter to the templates. the director now configures the L3 Agent correctly.

BZ#1244913

Due to a casting bug, 'L3HA' was passed as 'True' for overcloud Neutron configuration when deploying with only one controller. Because Neutron requires a minimum of 2 L3 agents to ensure native L3 HA, it complained when the overcloud tenant networking was set up after an otherwise successful deployment. As a consequence, after a successful deployment and when creating tenant Neutron networks or routers, the following error was returned from Neutron:

```
Not enough l3 agents available to ensure HA. Minimum required 2, available 1.
```

This bug has been fixed by correctly casting the number of controllers to an integer before checking for enablement of L3 HA. As a result, deployments with one controller have 'L3HA' as 'False' in the `/etc/neutron/neutron.conf` file on the controller node. Deployments with multiple controllers will instead have this set to 'True'.

6.2.12. python-tuskarclient

BZ#1228433

python-tuskarclient had no support to specify a custom CA certificate to verify SSL certificates from the Tuskar API server. Connections failed due the server using an certificate set unknown to the client. This fix adds the `--os-cacert` option to python-tuskarclient, which allows specification of the CA certificate path. This provides successful communication with the API server.

6.2.13. rhel-osp-director

BZ#1225016

On the Overcloud, the `glance_store` section in `/etc/glance/glance-api.conf` file sets `stores=glance.store.filesystem.Store`. This caused problems with image uploads due to the different store types. This fix modifies glance configuration to use `glance.store.http.Store` for the `stores` parameter and to include a backend for the store type to use.

BZ#1225621

Incorrect ordering of DHCP options the configuration file caused machines to fail on boot. This fix uses tags to the DHCP option to provide the correct ordering. Machine now chainload the boot with the iPXE ROM and then invoking the HTTP URL to continue the boot. This results in a successful boot.

BZ#1226097

The grub configuration set the kernel parameters to redirect the console to a serial port that might not be present. As a result, the node failed to boot. This fix disables console redirection to the serial port by default. The node now boots successfully.

BZ#1227421

Upstream Nova defaults to appending "novalocal" to host names. This causes problems with our tooling. The resulting nova host names don't correspond to the names expected by Puppet. This caused `CREATE_FAILED` on `ControllerNodesPostDeployment` in the Heat stack during Overcloud creation. This fix overrides the upstream default and does not set the "novalocal" suffix. Now post-deployment Puppet configuration reports a `CREATE_COMPLETE` on the Overcloud Heat stack.

BZ#1227940

The `deployment_mode` configuration for the Undercloud that could be set in `instack.answers` in tech preview has been removed. `deployment_mode` allowed choosing a "scale" mode that deployed individual roles onto specific nodes. This functionality has been replaced with the node tagging features part of the Automated Health Check (AHC) tools.

BZ#1228419

The director provides a method to isolate different network services on individual subnets and VLANs. These instructions are contained within the Basic and Advanced Overcloud Scenarios in the Red Hat Enterprise Linux OpenStack Platform 7.0 Director Installation and Usage guide.

BZ#1229372

Mixing DHCP and static routes both appear in the route table. However, the DHCP route had a metric of 100. This meant the static route with a metric of 0 was always used instead. As a workaround for using multiple DHCP links, you can configure one or more to ignore the default route offered by the DHCP server. Use the "defroute: false" declaration to accomplish this.

BZ#1230840

Previously, deployment specific values were not provided in OpenStack Integration Test Suite (tempest), resulting in the failure of some tempest tests.

With this update, a '--deployer-input' flag is added for the 'openstack overcloud validate' command so the administrator can provide a file (tempest.conf) containing the deployment specific values. As a result of using the '--deployer-input filename' flag, fewer tests result in failure.

BZ#1230966

Redis needs to use a separate VIP. When deploying with network isolation, the director automatically places Redis VIP on the Internal API VIP by default. Operators do have the ability to move Redis to another network using the ServiceNetMap parameter.

BZ#1233860

The director previously used a file with network and floating IP options for deployment. However, these network options were ignored in early versions of the director. This fix replaces the file with a set of command line arguments for "openstack overcloud director". The CLI tool now contains the necessary options for configuring network and floating IPs in the Overcloud.

BZ#1234856

Previously, an error in the name of the 'management_plan.uuid' variable led to attribute errors that would cause all deployments of Tuskar plans to fail. With this update, this variable name has now been corrected.

BZ#1235476
-

When deploying an Overcloud, the director placed the Public VIP services on the Provisioning network's "ctlplane". This meant you could not reach these services from outside of the Overcloud. This fix patches the Heat templates to place the Public VIP on the External network.

BZ#1235624

When deploying an Overcloud, the director placed the Public VIP services on the Provisioning network's "ctlplane". This meant you could not reach the horizon and keystone services from outside of the Overcloud. This fix patches the Heat templates to place the Public VIP on the External network.

BZ#1235667

Previously, in cases when ironic-discoverd submitted a 'set_boot_device' request with the same boot device already configured on the node, the DRAC driver in OpenStack Bare Metal Provisioning (ironic) tried to commit a configuration job against BIOS without any change, resulting in a failed deployment.

With this update, the DRAC driver in OpenStack Bare Metal Provisioning ignores requests where the target boot device is the same as the current one. As a result, the deployment succeeds.

BZ#1235908

Previously, the keystone token expired while Orchestration worked to deploy the Overcloud, resulting in the deployment failing with an authentication error.

With this release, the token timeout has been increased, resulting in the successful deployment of the Overcloud.

BZ#1236251

The default route was not set on the External network. This meant you could only access Horizon and the Public API from the same subnet as the Controller. This fix updates the Heat templates to include the ExternalInterfaceDefaultRoute parameter. This ensures a default route is available on the External interface.

BZ#1236578

The NeutronScale resource renamed neutron agents on Controller nodes. This caused an inconsistency with the "neutron agent-list" and as result Neutron reported errors of not having enough L3 agents for L3 HA. This fix removes the NeutronScale resource from Overcloud Heat templates and plans. NeutronScale does not appear in "neutron agent-list" and Neutron reports no errors.

BZ#1237064

Previously, the undercloud admin user did not have a 'swiftoperator' role on the service tenant. As a result, it was not possible for CloudForms Management Engine (CFME), which uses the admin user to access the swift objects.

With this update, the undercloud admin user is given the 'swiftoperator' role on the service tenant, thus, allowing the admin user access to the swift objects by specifying the service tenant in the API request.

BZ#1237144

The NeutronScale resource renamed neutron agents on Controller nodes. This caused an inconsistency with the "neutron agent-list" and as result Neutron reported errors of not having enough L3 agents for L3 HA. This fix removes the NeutronScale resource from Overcloud Heat templates and plans. NeutronScale does not appear in "neutron agent-list" and Neutron reports no errors.

BZ#1237145

With this release, a new enhancement adds NFS backend for the Block Storage service to provide a greater variety of possible Block Storage backends.

As a result, Overcloud Block Storage service can be configured with a NFS backend with the following parameters:

- * CinderEnableNfsBackend
- * CinderNfsMountOptions
- * CinderNfsServers

BZ#1237150

The glance backend was hard coded to use swift. This meant you could not use other file system types, such as NFS, for glance. This release adds a new enhancement that provides a NFS backend for glance. You can now configure the Overcloud's glance service an NFS backend using the following parameters:

- * GlanceBackend
- * GlanceFilePcmkManage
- * GlanceFilePcmkDevice
- * GlanceFilePcmkOptions

BZ#1237235

The OpenStack Dashboard (Horizon) is now a part of the high availability architecture. The Dashboard is now a resource that Pacemaker manages.

BZ#1237329

Pacemaker and ironic fought for control over power management, which caused issues with fencing. This fix sets `force_power_state_during_sync=False` in `/etc/ironic/ironic.conf` by default. This stops ironic automatically restoring the power state of the node during its synchronization. Pacemaker can now successfully fence the node.

BZ#1238217

The lack of a CLI parameter to set `NeutronExternalNetworkBridge` caused problems when assigning floating IPs. This means the only way to set this parameter is through custom environment file for network isolation. For example:

```
parameter_defaults:
  # Set to "br-ex" if External is on native VLAN
  NeutronExternalNetworkBridge: ""
```

Set this parameters to `'` if the floating IP network is on a VLAN and to `'br-ex'` if on a native VLAN on the br-ex bridge. This configuration allows the Neutron bridge mapping to work correctly for the environment. This is documented in the Red Hat Enterprise Linux OpenStack Platform 7 Director Installation and Usage guide

BZ#1238434

The `ipxe-bootimgs` package was unavailable in the default Red Hat Enterprise Linux repository. It was only available in the Red Hat Enterprise Linux - Optional repository. This caused deployments without the Optional channel to fail. This fix adds this package to the director channel. Deployments now work without the Optional channel enabled.

BZ#1238750

The `NeutronScale` resource renamed neutron agents on Controller nodes. This caused an inconsistency with the `"neutron agent-list"` and as result Neutron reported errors of not having enough L3 agents for L3 HA. This fix removes the `NeutronScale` resource from Overcloud Heat templates and plans. `NeutronScale` does not appear in `"neutron agent-list"` and Neutron reports no errors.

BZ#1238844

The Overcloud's configured its Heat component incorrectly and lacked settings for `heat_stack_user_role`, `stack_domain_admin`, and `stack_domain_admin_password`. This fix correctly sets the user and admin roles in `/etc/heat/heat.conf`.

BZ#1238862

Deployed Overclouds configured the Heat CloudFormation API to use an

auth_url pointing at localhost. However, Keystone does not listen on localhost. This caused an unusable Heat CloudFormation API. This fix changes the auth_url option in /etc/heat/heat.conf to the IP address where Keystone is listening on the Internal API network. The Heat CloudFormation API now functions correctly.

BZ#1240449

The Overcloud configured the heat service with instance_user=heat-admin. This meant SSH communication into heat-provisioned guest VMs required the heat-admin user. This fix sets instance_user to an empty value. Now you can SSH into guest VMs using the default image user.

BZ#1240824

The number of connections to the database scaled depending on the number of Controllers and the cores of each Controller. In HA scenarios with three controllers where each has more than 12 cores, the database connections could reach the default max_connections value set to 1024. This caused services to not respond to requests. As a workaround, increase the max_connections limit with the following command:

```
$ openstack management plan set [tuskar_plan_uuid] -P "Controller-1::MySQLMaxConnections=4096"
```

Replace [tuskar_plan_uuid] with the actual plan UUID, which you can find with:

```
$ openstack management plan list
```

To increase the max_connections value when deploying with the --templates argument, provide to the deploy command an additional customization environment file containing the following:

parameters:

```
MySQLMaxConnections: 4096
```

Add it to the deploy command with:

```
$ openstack management deploy --plan overcloud -e /path/to/custom_environment_file.yaml
```

BZ#1241131

Nodes would end up out of synchronization due to a lack of access to NTP servers. This was because not all nodes routed access to the required servers (NTP, DNS, etc). This fix sets the Undercloud as a gateway for non-Controller nodes. This provides non-Controller nodes with access to external services such as DNS and NTP, which aids synchronization.

BZ#1241422

SELinux was set to enforcing mode on Ceph OSD nodes. However, according

to official Ceph documentation, SELinux should be set to permissive mode on Ceph OSD nodes. This fix sets SELinux to permissive on Ceph OSD nodes.

BZ#1242052

The timeout for Pacemaker service start-up was 20 seconds. Sometimes start-up exceeded this time limit and caused hung deployments. This fix increases the timeout to 60 second. Pacemaker services now start correctly and the deployment completes.

6.3. RHSA-2015:1862 — MODERATE: RED HAT ENTERPRISE LINUX OPENSTACK PLATFORM 7 DIRECTOR UPDATE

The bugs contained in this section are addressed by advisory RHSA-2015:1862. Further information about this advisory is available at <https://access.redhat.com/errata/RHSA-2015:1862.html>.

6.3.1. ahc-tools

BZ#1245212

SSL configuration on the director caused the Automated Health Check (AHC) tools to fail due to not using internal endpoints for certain components. This fix changes the configuration to use internal endpoints. The AHC tools now run without SSL errors.

6.3.2. instack-undercloud

BZ#1223022

A missing firewall rule restricted access to the Ceilometer API. This fix adds the firewall rule. Users now have access to the Ceilometer API.

BZ#1226376

The director's iptables previously denied port 9696. This rejected all requests to the Neutron API except for those coming from localhost. This fix adds an iptables rule to accept TCP traffic for port 9696. Remote connections now have access to the Neutron API.

BZ#1236707

`undercloud_heat_encryption_key` was incorrectly documented in `undercloud.conf.sample` as accepting values of size 8, 16, or 32 characters. Only values of size 16, 24, or 32 characters are actually accepted. If a value that had the non-accepted size was used, the Undercloud configuration script would fail with an error similar to:

```
Error: 8 is not a correct size for auth_encryption_key parameter, it must be either 16, 24, 32 bytes long. at
```

`/etc/puppet/modules/heat/manifests/engine.pp:106` on node `undercloud.local.dev`

This fix updates the `undercloud.conf.sample` with the correct documentation to indicate the accepted sizes of this parameter.

BZ#1243121

The Overcloud deployment used a default port quota of 50 for Neutron networking, which caused failures in larger deployments. This fix disables the port quota. Larger Overcloud deployments no longer fail from lack of Neutron ports.

BZ#1247015

The Undercloud configuration script ignored rabbit user details in `undercloud.conf` and did not create the necessary user for rabbitmq. This caused an incorrect rabbitmq configuration that resulted in a failed Undercloud configuration. This fix adds code to the undercloud configuration script that creates the requested rabbitmq user with the requested password. All services now connect to rabbitmq using the requested username and password.

BZ#1251566

The director's database (MariaDB) only accept a maximum of 1024 connections. An Undercloud with a high number of CPU cores (typically 24 or more) exhausted these database connections due to the number of OpenStack API workers spawned. This fix configures the Undercloud to accept 4096 connections for MariaDB. All services now connect to MariaDB when needed.

BZ#1256477

Nodes registered with an unresponsive IPMI IP address caused the sync power state periodic task to hang for a default 10 minutes. This resulted in unresponsive behavior from Ironic. This fix lowers the default IPMI retry timeout. Now unresponsive nodes report failures faster and do not hang on the sync power state periodic task.

6.3.3. openstack-ironic-discoverd

BZ#1252437

The inspection process picked a random root disk to report as `local_gb`. This often returned the wrong `local_gb` value, which would differ from run to run on machines with multiple hard disks. This fix sorts the order of the disks before picking the first one. The inspection process now provides a consistent `local_gb` value.

6.3.4. openstack-tripleo

BZ#1243472

Updating an Overcloud set the UpdateIdentifier parameters for each node in the director's Overcloud plan. However, deleting the Overcloud stack and redeploying it resulted in failure if the UpdateIdentifier parameters were set due to no preset repositories upon deployment. This fix stops the Overcloud from setting the UpdateIdentifier parameter on each node. This results in a successful Overcloud deployment.

BZ#1252509

The "openstack overcloud update stack -i" command did not handle invalid regular expressions. This caused command termination with an error:

```
openstack unexpected end of regular expression
```

This fix detects invalid regular expressions and provides a warning to the user about the invalid value. The command now handles invalid expressions and prompts the user to enter a new value.

6.3.5. openstack-tripleo-heat-templates

BZ#1230844

This enhancement adds support for the Nexus-9k ML2 Neutron plugin. This includes environment configuration in the TripleO Heat Template collection as well as configuration in the Openstack Puppet Module collection.

BZ#1230850

This enhancement adds support for the Cisco UCSM Neutron ML2 plugin. This includes environment configuration in the TripleO Heat Template collection as well as configuration in the Openstack Puppet Module collection.

BZ#1233949

Load balancing for httpd was incorrectly configured on the Overcloud, which meant VIPs were not used when accessing Horizon. This fix properly enables the load balancing for httpd. Now Horizon is accessible through VIPs.

BZ#1236136

All keystone endpoints are on the External VIP. This means all API calls to keystone happen over the External VIP. There is no workaround at this time.

BZ#1249832

This enhancement increases the levels of configuration for the Overcloud's Neutron service. Customers can now configure values for `core_plugin`, `type_drivers`, and `service_plugins` through the director.

BZ#1252219

The bonded NIC templates use a specific parameter to name the main bridge that connects all VLANs on Controller and Compute nodes. The Overcloud's networking expects the same bridge name on both Controller and Compute nodes. However, the default bridge name is different for each node type (`br-ex` for Controller, `br-bond` for Compute). This results in missing packets on the bonded interface and a faulty Overcloud networking configuration. This fix removes the hardcoded value for the bond name in the Compute node NIC configuration. Using the input value for `bridge_name` instead ensures the Controller and Compute nodes have the same bridge name (defaults to `"br-ex"`).

BZ#1254897

Overcloud deployments did not consume nor affect the Neutron mechanism driver parameter passed as hiera data from the corresponding Heat template. This meant Overcloud deployments contained an unexpected Neutron configuration, with both `openvswitch` and `linuxbridge` configured as `mechanism_drivers` in `/etc/neutron/plugins/ml2/ml2_conf.ini` as such:

```
mechanism_drivers = openvswitch,linuxbridge
```

This fix ensures the Overcloud deployment correctly consumes the `neutron_mechanism_drivers` hiera data item passed from the Heat templates and sets this in the Neutron ml2 configuration on the controller node. You can also specify the `NeutronMechanismDrivers` Heat template parameter as a custom parameter and expect the corresponding ml2 configuration for Neutron.

BZ#1257414

Missing constraints between Pacemaker resources caused issues when starting or stopping the Controller cluster. This fix adds these constraints. Pacemaker resources now have the necessary relationships to function.

BZ#1262995

The default external network port was set to control plane VIP. This caused network validation to fail when running without network isolation enabled. This fix returns the correct default port when network isolation is not in use. The network validation now succeeds.

BZ#1265013

The collection of hostname to MAC mappings for Cisco ML2 Nexus support

The collection of hostname to IP mappings for Cisco IRE nodes support used the nameserver, which might be unavailable due to timing issues. This caused Overcloud configuration issues. This fix check if "hostname -f" fails. If it does, director appends the hostname explicitly with ".localdomain". Now the hostname look-up works regardless of nameserver timing issues.

6.3.6. openstack-tripleo-puppet-elements

BZ#1255423

Incorrect data type handling between Heat and Puppet caused complex parameters, such as JSON hashes, to pass incorrectly from Heat to Puppet. This fix improves the data types handling in the component which writes Hiera data from values received from Heat. Now passing JSON hashes from Heat to Puppet functions correctly.

6.3.7. openstack-tuskar

BZ#1253628

Support for using a pre-existing external Ceph Storage cluster with Overcloud deployments caused all deployments using a Tuskar plan to fail. This included deployments through the web UI and CLI when `--plan` is used instead of `--templates`. This was irrespective of whether any Ceph Storage nodes were deployed (external or otherwise). Failed deployments resulted in the following error:

```
ERROR: openstack ERROR: Property error : CephClusterConfig:
ceph_storage_count The Parameter (Ceph-Storage-1::CephStorageCount) was
not provided.
```

This fix modifies Tuskar so that it can work with nested references to the top level Heat templates scaling parameters (like `CephStorageCount` in this case). Deployments using Tuskar now function without encountering the `CephStorageCount` error.

6.3.8. os-cloud-config

BZ#1233564

This fix adds support for Cisco UCS machines to Ironic's power management control in the director. Cisco UCS nodes are manageable using the IPMI protocol, but some customers might want to use the specific Cisco UCS driver to manage more advanced features. Now the director supports power management for Cisco UCS machines.

BZ#1259393

This enhancement adds support for the `fake_pxe` Ironic driver for registering machines without power management to the director. Use the

fake_pxe driver as a fallback driver for machines without a power management system. Perform all power operations manually when using this driver.

BZ#1262454

The director expected power management authentication details when using the fake_pxe driver. This caused failure when registering nodes. This fix updates the os-cloud-config tool to disregard pm_addr, pm_password, and pm_user when using the fake_pxe driver. The director now successfully registers nodes using the fake_pxe driver.

6.3.9. python-hardware

BZ#1257517

In the discovery ramdisk, the python-hardware module ran a utility that returned a string with invalid UTF-8 encoding. This caused the hardware-detect command to exit with an error. As a result, the discovery ramdisk dropped to a dracut shell. This fix modifies the module to mitigate and resolve the error, rather than exit from the hardware-detect command. The ramdisk no longer drops to a dracut shell in this situation.

6.3.10. python-proliantutils

BZ#1248172

The iLO drivers lacked support for introspection cleaning tasks, which caused the ironic introspection to fail. A fix has been pushed to the proliantutils library to correct the iLO drivers. Alternatively, as a workaround, disable the cleaning task in /etc/ironic/ironic.conf:

```
[ilo]
...
clean_priority_reset_ilo=0
clean_priority_reset_bios_to_default=0
clean_priority_reset_secure_boot_keys_to_default=0
clean_priority_clear_secure_boot_keys=0
clean_priority_reset_ilo_credential=0
```

6.3.11. python-rdomanager-oscplugin

BZ#1231777

The "openstack overcloud deploy" command did not check available nodes for deployment. This caused failed deployments due if there were not enough nodes available. This fix adds a pre-deployment check to the CLI and checks the number of available nodes before creating or updating the Overcloud stack. Now if not enough nodes are available, users get an error message before Heat creates or updates the stack.

BZ#1235325

The "openstack baremetal configure boot" command attempted to configure nodes in maintenance mode. This caused boot configuration to fail. This fix skips nodes in maintenance mode. Now the boot configuration passes without error.

BZ#1241199

Running "openstack baremetal configure boot" overwrote the "capabilities" property of bare metal nodes. This removed the profile information for existing nodes. This fix changes the overwrite method to an append method. This no longer removes the profile information.

BZ#1243828

Network setup details were not passed to Tempest, which caused tests to fail when running "openstack overcloud validate". This fix generates deployer input at the end of deployment. However, you must now run Tempest manually. This fix also removes the "openstack overcloud validate" command from the director.

BZ#1243829

The "openstack overcloud image upload" command uploaded Overcloud images even if old versions existed in Glance. This resulted in images with duplicate names, which caused Overcloud creation to fail. This fix modifies the tool to skip existing images. The tool also includes an "--update-existing" option to update existing images. Overcloud creation now uses the new Overcloud images stored in Glance without failure.

BZ#1244001

Bulk introspection applied to all nodes including active nodes. However, bulk introspection failed on active nodes. This fix no longer applies introspection to active nodes.

BZ#1244856

Unclear help text and bug in parameter passing for the "openstack overcloud update stack" operation meant users needed to specify both plan and stack names. This fix removes the need for the stack name. However, the update command now requires either the Tuskar plan ID or the Heat template collection location.

BZ#1249640

The director generated no deployer input for Tempest after deployment. This caused possible missing tempest configurations options which were not auto-detected. This fix generates deployer input at the end of

deployment. However, you must now run Tempest manually. This fix also removes the overcloud validate command from the director.

BZ#1253777

The "--ntp-server" option was not provided for some HA Overcloud deployments. This caused the clock on Controller nodes to drift, which caused problems in Keystone. This fix sets the "--ntp-server" option to mandatory for deployments with multiple Controller nodes. Clocks on Controller nodes are now synchronized.

BZ#1265010

Overcloud updates passed default environment files to Heat. If the Overcloud creation used additional environment files that were not passed during the update, the Overcloud would update the resource registry definitions as per the default environment files, which in turn deleted some of the Overcloud's Heat resources. Now, the default environment files are no longer sent to Heat on update, and Heat does not delete the Overcloud resources.

6.3.12. rhel-osp-director

BZ#1229811

This enhancement adds support for the Cisco N1kV plugin. This includes environment configuration in the TripleO Heat Template collection.

BZ#1241720

This enhancement adds support for the Cisco N1kV VEM module. This includes environment configuration in the TripleO Heat Template collection.

BZ#1255910

When deleting a node in the Overcloud, the Heat stack's ComputeCount parameter calculated the number of nodes. However, Heat did not update parameters if a scale up operation failed. This meant the number of nodes that Heat returned in parameters did not reflect the real number of nodes. This caused problems with the number of nodes deleted on a failed stack. This fix ensures Heat updates the parameters even if a scale operation failed previously. Now the director deletes the requested nodes when running "overcloud node delete" on a stack where scale up operation failed before.

BZ#1265777

When using static IPs on the ctlpane network, DNS must be configured on the Overcloud nodes. Previously, DHCP provided the configured DNS servers for the Overcloud nodes.

To configure DNS when using static IPs set the new `DnsServers` parameter and include it in the Heat environment like so:

```
parameter_defaults:
  DnsServers:
    - <dns server ip address>
    - <dns server ip address 2>
```

Use IP addresses for the DNS servers to use. You can only specify either one or two DNS servers.

6.3.13. vulnerability

BZ#1261697

A flaw was discovered in the pipeline ordering of OpenStack Object Storage's `staticweb` middleware in the `swiftproxy` configuration generated from the `openstack-tripleo-heat-templates` package (OpenStack director). The `staticweb` middleware was incorrectly configured before the Identity Service, and under some conditions an attacker could use this flaw to gain unauthenticated access to private data.

6.4. RHSA-2015:2650 — MODERATE: RED HAT ENTERPRISE LINUX OPENSTACK PLATFORM 7 DIRECTOR UPDATE

The bugs contained in this section are addressed by advisory RHSA-2015:2650. Further information about this advisory is available at <https://access.redhat.com/errata/RHSA-2015:2650.html>.

6.4.1. openstack-tripleo-heat-templates

BZ#1241434

Heat templates lacked the `*RemovalPolicies` parameters. This meant it was not possible to delete specific nodes when using Heat templates directly (i.e. not through Tuskar). This update adds the `*RemovalPolicies` parameters. Now a user can specify particular nodes to remove by setting the `*RemovalPolicies` parameters.

BZ#1285079

Previously, orphaned OpenStack Networking L3 agent `keepalived` processes were left running by OpenStack Networking's `netns-cleanup` script. As a result, the OpenStack Networking tenant router failover did not work during the controller node update in the overcloud.

With this update, `keepalived` processes are cleaned up properly during the controller node update. As a result, OpenStack Networking tenant router failover works normally and the high availability of the tenant network is preserved.

BZ#1290582

Previously, during the initial Overcloud deployment, there existed a race condition between the puppet trying to stop the neutron-server and the Pacemaker trying to start the neutron-server. The neutron-server would often be left stopped on the Overcloud controllers, even though the deployment indicated it was successful. This was because the request to stop neutron-server eventually succeeded, although it would be not reported to Orchestration.

With this update, the puppet manifest is fixed to only conditionally stop the neutron-server if the Pacemaker is not already managing the neutron-server resource. As a result, the initial deployments succeed and the neutron-server is running in the Overcloud.

6.4.2. python-rdomanager-oscplugin

BZ#1245737

Previously, hard-coded parameters were being passed directly to Orchestration. As a result, the parameters could not be overridden properly.

With this update, a custom environment file from the parameters collected is generated and pass as 'parameter_defaults', allowing parameters to be overridden.

BZ#1259084

Previously, the 'debug' parameter was enabled and hard-coded in the overcloud deployment code and there was no way to disable the debug mode by an user.

With this update, the 'debug' parameter is removed from the default hard-coded parameters in the overcloud deployment code. As a result, the user can now control debug mode in the environment file used to deploy the overcloud.

BZ#1260776

This update removes an incorrect warning when deploying the Red Hat Enterprise Linux OpenStack Platform director.

BZ#1261863

Previously, while checking the node configuration deployment validation checked for all Ironic nodes, including those in the maintenance mode.

With this update, the nodes in maintenance mode are skipped by the validation step as they cannot be deployed.

BZ#1265714

Previously, the 'tempest-deployer-input.conf' file contained incorrect `stack_owner_role`. So using the 'tempest-deployer-input.conf' file for post-install validation caused more Tempest test failures.

With this update, the value in the 'tempest-deployer-input.conf' file generated during deployment is changed. As a result, less number of Tempest tests will fail during post-install validation.

BZ#1267558

Previously, breakpoints were not removed when an update operation failed. If a user ran the "openstack overcloud update" command and it failed, it is possible that the subsequent stack-update command (e.g. "openstack overcloud deploy") might be stuck in the 'IN_PROGRESS' state waiting for the removal of breakpoints.

With this update, all existing CLI commands explicitly remove any existing breakpoints when running a stack-update operation. As a result, the stack-update operations do not get stuck in the 'IN_PROGRESS' state while waiting for the breakpoint removal.

BZ#1267855

Previously, the base resource registry environment was included for all overcloud stack updates, which meant customizations may be lost unless all environment files are repeated in order when calling "openstack overcloud deploy".

With this update, it is possible to call "openstack overcloud deploy" with no environments without losing customizations. If any environment files are specified, then all environment files must be specified again in the desired order.

BZ#1268415

Previously the base resource registry environment was included for all overcloud stack updates, which meant customizations may be lost unless all environment files are repeated in order when calling "openstack overcloud deploy".

With this update, it is possible to call "openstack overcloud deploy" with no environments without losing customizations. If any environment files are specified, then all environment files must be specified again in the desired order.

BZ#1290796

Previously, when scaling out the Compute nodes in the Overcloud after an update was performed, the default `UpdateIdentifier` parameter in the Heat stack caused the new Compute node to attempt an update as soon as it was coming up. Since the yum repositories were not configured on the new

Compute nodes yet, it would cause the update to fail, which in turn caused the scale out to fail.

With this update, the client, `python-rdomanager-oscpugin`, does not clear the `UpdateIdentifier` parameter on the subsequent stack-update attempts (including the scale out) until after the initial update has been completed. As a result, scale out attempts after the update now succeeds.

6.4.3. rhel-osp-director

BZ#1266910

Previously, a Pacemaker constraint between the `neutron-server` and the `neutron-ovs-cleanup` processes meant that the over-cleanup restarted whenever the `neutron-server` did. The cleanup by the `ovs-cleanup` (and associated `netns-cleanup`) should only occur when a node was being evacuated or rebooted and not for a `neutron-server` restart. As a result of this constraint, the `neutron-server` needed long startup to rebuild the data plane and ultimately caused issues for the dependent DHCP and L3 agents.

With this update, the constraint between the `neutron-server` and the `ovs-cleanup/netns-cleanup` was removed. This means that the `ovs-cleanup/netns-cleanup` does not re-run after `neutron-server` is restarted (for example, because `haproxy` was). The result is the two constraints chains for OpenStack Networking: `ovs-cleanup-->netns-cleanup-->openvswitch-agent`, and then `neutron-server-->openvswitch-agent-->dhcp-agent-->l3-agent-->metadata-agent`. As a result, when `haproxy` is restarted or when `neutron-server` is, more specifically, the `ovs` and `netns` cleanup scripts will not be re-run, so the OpenStack Networking service startup will proceed normally.

BZ#1272347

With this update, the default network where the `'KeystoneAdminVip'` is placed was changed from `'InternalApi'` to `'ctlplane'` so that the post-deployment Identity service initialization step could be carried by the Undercloud over the `'ctlplane'` network. Relocating the `'KeystoneAdminVip'` causes a cascading restart of the services pointing to the old `'KeystoneAdminVip'`.

As a workaround to make sure the `KeystoneAdminVip` remains on the `'InternalApi'` network, a customized `'ServiceNetMap'` must be provided as deployment parameter when launching an update from the 7.0 release. A sample Orchestration environment file passing a customized `'ServiceNetMap'` is as follows:

```
parameters:
  ServiceNetMap:
    NeutronTenantNetwork: tenant
    CeilometerApiNetwork: internal_api
    MongoDBNetwork: internal_api
```

```

CinderApiNetwork: internal_api
CinderIscsiNetwork: storage
GlanceApiNetwork: storage
GlanceRegistryNetwork: internal_api
KeystoneAdminApiNetwork: internal_api
KeystonePublicApiNetwork: internal_api
NeutronApiNetwork: internal_api
HeatApiNetwork: internal_api
NovaApiNetwork: internal_api
NovaMetadataNetwork: internal_api
NovaVncProxyNetwork: internal_api
SwiftMgmtNetwork: storage_mgmt
SwiftProxyNetwork: storage
HorizonNetwork: internal_api
MemcachedNetwork: internal_api
RabbitMqNetwork: internal_api
RedisNetwork: internal_api
MysqlNetwork: internal_api
CephClusterNetwork: storage_mgmt
CephPublicNetwork: storage
ControllerHostnameResolveNetwork: internal_api
ComputeHostnameResolveNetwork: internal_api
BlockStorageHostnameResolveNetwork: internal_api
ObjectStorageHostnameResolveNetwork: internal_api
CephStorageHostnameResolveNetwork: storage

```

If any additional binding network from the above has been customized then that setting has to be preserved as well.

As a result of the workaround changes, the 'KeystoneAdminVip' is not relocated on the 'ctlplane' network so that no services restart needs to be triggered.

6.4.4. vulnerability

BZ#1272297

It was discovered that Director's NeutronMetadataProxySharedSecret parameter remained specified at the default value of 'unset'. This value is used by OpenStack Networking to sign instance headers; if unchanged, an attacker knowing the shared secret could use this flaw to spoof OpenStack Networking metadata requests.

BZ#1281777

A flaw was found in the director (openstack-tripleo-heat-templates) where the RabbitMQ credentials defaulted to guest/guest and supplied values in the configuration were not used. As a result, all deployed overclouds used the same credentials (guest/guest). A remote non-authenticated attacker could use this flaw to access RabbitMQ services in the deployed cloud.

6.5. RHBA-2015:2651 — RED HAT ENTERPRISE LINUX OSP 7 DIRECTOR BUG FIX ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2015:2651. Further information about this advisory is available at <https://access.redhat.com/errata/RHBA-2015:2651.html>.

6.5.1. instack-undercloud

BZ#1272509

A missing dnsmasq lease file caused "instack-virt-setup" to fail when retrieving the host IP. This fix retrieves the IP address from a different source and "instack-virt-setup" now succeeds.

6.5.2. openstack-tripleo

BZ#1259076

When updating an Overcloud, the director parsed each breakpoint individually. This added to the amount of time it took to update the Overcloud. This fix parses the whole stack once for all breakpoints, reducing the amount of time to update.

BZ#1266102

The director updates Heat parameters even if a stack-update operation failed. The Heat stack *Count parameters (e.g. ComputeCount) might not reflect the real number of nodes in an Overcloud. This caused the "overcloud node delete" command to delete an incorrect number of nodes. This fix modifies the "overcloud node delete" command to compute the current node count from the real number of servers in ResourceGroup instead of using stack parameters. Now the director deletes the correct number of nodes.

BZ#1286774

The timeout for Overcloud updates was set to 90 minutes. This caused incomplete updates that took longer than 90 minutes. This fix extends the timeout to 240 minutes. Longer updates now complete successfully.

6.5.3. openstack-tripleo-heat-templates

BZ#1241434

Heat templates lacked the *RemovalPolicies parameters. This meant it was not possible to delete specific nodes when using Heat templates directly (i.e. not through Tuskar). This update adds the *RemovalPolicies parameters. Now a user can specify particular nodes to remove by setting the *RemovalPolicies parameters.

6.5.4. os-apply-config

BZ#1271687

The os-apply-config tool included unit tests as part of its installation. These tests were meant to be omitted. This fix updates the spec file to remove these tests. Unit tests are no longer included with os-apply-config.

6.5.5. os-cloud-config

BZ#1258651

When registering new nodes and using the existing JSON file to store their info, the director would overwrite the capabilities of existing nodes. This fix skips the existing nodes. Now users can add new nodes to the original JSON file without erasing capabilities for existing nodes.

BZ#1274241

This enhancement adds support for Fujitsu's iRMC Ironic driver in the director. The director now controls the power management of iRMC nodes in the Overcloud.

6.5.6. os-collect-config

BZ#1271700

The base RPMs included unit tests. This meant the director had unnecessary files that consumed extra space. This fix removes the tests from the base RPMs. Unit tests are now only shipped with the source RPMs.

BZ#1272254

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.

6.5.7. os-net-config

BZ#1244010

This enhancement adds Linux bonding configuration through the director. The director used only OVS and VLAN bonding previously. Linux bonding provides increased performance and additional bonding modes.

6.5.8. rhel-osp-director

BZ#1173970

No regular database maintenance process existed in previous versions of the director. As a result, the director's database grew without limit. This fix adds a cronjob to flush expired tokens from the database. This cleans the database periodically and reduces its size.

BZ#1244946

An issue with Red Hat Enterprise Linux 7 boot settings caused unpredictable NIC naming on Overcloud images. This fix updates the Overcloud images to use Red Hat Enterprise Linux 7.2, which contains a fix for this issue.

BZ#1247358

In rare cases, RabbitMQ fails to start on deployment. As a workaround, manually start RabbitMQ on nodes:

```
[stack@director ~]$ ssh heat-admin@192.168.0.20  
[heat-admin@overcloud-controller-0 ~]$ pcs resource debug-start rabbitmq
```

Then rerun the deployment command on the director. The deployment now succeeds.

BZ#1272176

This enhancement upgrades the Overcloud image content to Red Hat Enterprise Linux 7.2 content, including the latest version of Pacemaker. The previous Overcloud image used Red Hat Enterprise Linux 7.1 content.

BZ#1272302

Rerunning the director configuration script failed while restoring an Undercloud. This fix makes the configuration script re-runnable without requiring database clean-up. Users can now rerun the director configuration script without failure.

BZ#1275439

This feature allows the reapplication of Puppet manifests on a deployed Overcloud. This ensures the overcloud has the desired configuration or can recover accidentally amended or deleted configuration files.

To have Puppet run again on the Overcloud nodes, omit the "--templates" option but include the following two environments files at the beginning of your deployment:

```
* /usr/share/openstack-tripleo-heat-templates/overcloud-resource-registry-puppet.yaml
```

```
* /usr/share/openstack-tripleo-heat-templates/extraconfig/pre_deploy/rhel-registration/rhel-registration-resource-registry.yaml
```

For example:

```
$ openstack overcloud deploy -e ~/templates/overcloud-resource-registry-puppet.yaml -e ~/templates/extraconfig/pre_deploy/rhel-registration/rhel-registration-resource-registry.yaml [additional arguments from initial deployment]
```

BZ#1275814

Old puppet manifests were reapplied during the update process when they should not have been. This had the potential to take the cluster services down in the Overcloud. The agent on the overcloud nodes caused the reapplication of the old Puppet manifests because their state was saved in tmpfs mounted directory under `/var/run/`. This directory is lost on reboot. This update moves the directory from `/var/run/heat-config/deployed` to `/var/lib/heat-config/deployed`, which allows the deployed state to persist across reboots.

BZ#1275986

Pacemaker prevented Puppet from restarting services during an Overcloud update. This caused Puppet runs to fail, which caused the Overcloud update to fail. This fix puts Pacemaker into maintenance mode for the duration of the Puppet run during the Overcloud update. Pacemaker now allows Puppet to restart services and now the Overcloud update does not fail anymore.

BZ#1278004

Configuration records for heat-config were stored under `/var/run`, which were lost on node reboot. The cause for this was reported and fix (see BZ#1278181) but had an impact on the OSP director; Puppet reapplied old manifests during an Overcloud update, which caused various Pacemaker cluster errors. The fix for BZ#1278181 moves the configuration records from `/var/run/heat-config` to `/var/lib/heat-config`. This fix also includes the "heat-config-rebuild-deployed" script to rebuild the configuration records. Now updates are possible without reapplying old Puppet manifests. Make sure to follow the documented update procedure, including running the "heat-config-rebuild-deployed" script on each node.

BZ#1278430

A missing SELinux rule caused Glance to return an "Invalid OpenStack Identity credentials" error. This fix adds the SELinux rule. Now Glance authenticates successfully.

BZ#1279649

■

Due to various issues with updating the Overcloud heat stack, updates from 7.0 to 7.1 with network isolation would not work. This fix corrects these issues. Now you can update deployments from 7.0 to 7.2 directly. Updates from 7.0 to 7.1 should be skipped. Ensure to follow the documented update procedure, as additional Heat environment files are needed depending on the configuration of particular deployments.

BZ#1279652

Orphaned OpenStack Networking L3 agent keepalived processes were left running by OpenStack Networking's "netns-cleanup" script. As a result, the OpenStack Networking tenant router failover did not work during the Controller node update in the Overcloud. This fix ensures the keepalived processes are cleaned up properly during the Controller node update. As a result, OpenStack Networking tenant router failover works normally and the high availability of the tenant network is preserved.

BZ#1283561

Strict timeouts on starting and stopping services through Pacemaker caused updates to fail. This fix increases the timeouts for starting and stopping services through Pacemaker. Now services start and stop within the timeout limit and the Overcloud now succeeds.

6.6. RHBA-2015:2652 — OPENSTACK-NEUTRON BUG FIX ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2015:2652. Further information about this advisory is available at <https://access.redhat.com/errata/RHBA-2015:2652.html>.

6.6.1. openstack-neutron

BZ#1253953

Previously, when HA routers were scheduled to multiple nodes, each such replica of the router had its own copy of its internal and external ports, however, from neutron's perspective each such port was bound only to a single host. With HA routers, only one replica of the router is active at any point in time, but the router's ports may be bound to a host that is in standby mode.

As a result, l2pop used the port binding information to configure flows. Since the neutron port for replicated interfaces could be bound to the wrong host, l2pop may have broken connectivity by configuring tunnel endpoints to the wrong host, or by configuring unicast openflow rules that point to a standby node. Additionally, some ML2 mechanism drivers would rely on the port binding information to configure ToR switches or other network gear, which was being misconfigured.

With this update, whenever keepalived performs a state transition, it notifies the L3 agent, which then notifies the neutron-server. The server then updates the port's binding information to point to the new active node. As a result, l2pop and other ML2 mechanism drivers now have a correct view of the external environment, with router ports owned by HA routers always being bound to the active node.

BZ#1256816

Previously, in certain circumstances (such as deployments using a vendor-specific implementation of the neutron L3 API), the neutron router was not available to provide the IP route for the metadata service.

This issue can be addressed using DHCP to allocate this information. Setting 'force_metadata = False' causes the DHCP server to append specific host routes to the DHCP request. As a result of performing this configuration change, the metadata service will be activated for all networks.

BZ#1268244

Prior to this update, the netns pacemaker OCF resource did not perform a full cleanup of the neutron netns services.

As a result, some of those services were orphaned, and were never restored by the l3-agent because they were seen as running, but were actually disconnected.

This update addresses this by adding the missing cleanup steps to the netns cleanup OCF resource.

BZ#1268859

Previously, metadata-proxy could not be spawned in the DHCP namespace if the network was attached to any router.

Consequently, a network could not be created if the router required a metadata-proxy process in the DHCP namespace.

This update resolves this issue by adding the new config option 'force_metadata' for dhcp_agent.ini. As a result, setting 'force_metadata' to 'True' will cause the metadata-proxy to always be spawned in the DHCP namespace, even if the network is attached to a router.

BZ#1269849

Prior to this update, the Linux iptables implementation of security groups included a default rule to drop any INVALID packets.

Consequently, it was possible that iptables could block legitimate traffic as INVALID, such as SCTP protocol.

This update addresses this issue by processing user-defined iptables rules before the INVALID DROP rule.

BZ#1274880

This neutron rebase package includes a number of notable enhancements and fixes under version 2015.1.2:

* Layer 3 High Availability:

- Fixed race condition when starting radvd processes for IPv6 networks
- Gratuitous APR updates are now repeated
- Fixed HA routers when l2population ML2 driver was used
- Fixed a bug where a HA router failed while configuring IPv6 Router

Advertisements on its external gateway

- It is now possible to configure the underlying physical network for VRRP traffic

*** L3:**

- Stale metadata processes are now cleaned up on sync
- Prevents attaching an interface to a router if the port does not have an IP address assigned
- Gratuitous ARPs are now skipped for IPv6 addresses

*** Distributed Virtual Routing (DVR):**

- Service port ARP is now broadcast
- Routers are now unscheduled if all remaining ports are not bound to the node

*** Security Groups:**

- Fixed ipset cleanup on last security group rule removal
- Fixed ipset cleanup if requested set does not exist
- IPtables manager is significantly optimized for performance
- Fixed interaction with LBaaS ports
- More fixes for default security group creation

*** DHCP:**

- Fixed a bug where some IPv6 addresses might miss name resolution settings
- Scheduler is optimized to guarantee the configured number of agents serving a network
- Fixed a bug where tunnels were not created on failover, when using the l2population ML2 driver

*** ML2 plugin:**

- Fixed rare race condition where a port and its network were removed in parallel

*** Open vSwitch (OVS):**

- Do not use ARP responder for IPv6 addresses

*** SR-IOV:**

- Fixed setting admin_state_up for ports

*** Linux Bridge:**

- Fixed race condition on bridge cleanup
- Tap device MTU is now set according to underlying physical device
- Added ARP spoofing protection support (disabled by default)

*** Port Security:**

- Fixed late enablement of the extension for existing networks

*** API:**

- Allow to unset description for an agent

BZ#1281432

Prior to this update, processing router information on L3 agent synchronization was performed inefficiently. Consequently, the neutron server load may have been unexpectedly high when using large numbers of

routers under non-extreme conditions.

This update addresses this issue by improving query efficiency, and removing unnecessary operations on synchronization.

As a result, neutron server CPU usage is greatly reduced when large numbers of routers are configured.

6.6.2. openstack-neutron-fwaas

BZ#1274889

This FWaaS rebase package includes a notable fix under version 2015.1.2

- Fixed DB tracebacks on multiple FWaaS API operations (rule insert, rule remove, and others)

6.6.3. openstack-neutron-lbaas

BZ#1274881

This LBaaS rebase package includes a number of notable enhancements and fixes under version 2015.1.2

- Gracefully error out when attempting to delete a port attached to a VIP
- device_id is now set for a LBaaS port on creation, to prevent nova from booting an instance using the port

6.6.4. openstack-neutron-vpnaas

BZ#1274891

This VPNaaS rebase package includes a notable fix under version 2015.1.2

- Confirms that the file containing the pre-shared key for VPN connections is not world-readable

6.7. RHBA-2016:0264 — RED HAT ENTERPRISE LINUX OSP 7 DIRECTOR BUG FIX ADVISORY

The bugs contained in this section are addressed by advisory RHBA-2016:0264. Further information about this advisory is available at <https://access.redhat.com/errata/RHBA-2016:0264.html>.

6.7.1. instack-undercloud

BZ#1286756

This firewall configuration for the Undercloud lacked certain ports, which resulted in dropped packets for Internal API messages. This fix adds the missing ports (13000, 13774, 13696, 13385, 13292, 13696, 13004, 13080, 13385) to the Undercloud's firewall rules. The Internal API now accepts messages on these ports.

BZ#1304441

The Undercloud's firewall lacked a port for Ceilometer's Public API over SSL. This fix adds the port (13777) to the Undercloud's installation script. Now Ceilometer accepts Public API requests over SSL.

BZ#1305918

The Undercloud installation script recreated users on subsequent runs. This causes the service user IDs to change, which causes trust issues for running certain services. This fix stops the installation script from recreating users. Now service user IDs remain consistent with their respective services.

6.7.2. ipxe**BZ#1300702**

This enhancement provides new iPXE images for the Red Hat OpenStack Platform director. The new iPXE images add extra network boot support for certain NICs

6.7.3. openstack-tripleo-common**BZ#1301763**

The 'openstack overcloud update' command sought a list of events for each resource. When listing events for a resource, the Heat API returned a HTTP 404 error (Not Found) for resources with no events. The resource was considered as non-existent (due to the 404 error) and the client would fail. This occurred in situations where a previous update would end after adding a resource to the stack but before any events occur, such as a resource waiting at a breakpoint at the time the update ended. This fix adds error handling to the client, which resolves issues with the 404 error.

6.7.4. openstack-tripleo-heat-templates**BZ#1238460**

Previous Overcloud images used EDT for the timezone, which caused problems for users outside of the EDT timezone. This fix adds a customizable 'TimeZone' parameter to the Heat template collection. Users can set the 'TimeZone' parameter to their own timezone. If blank, it defaults to UTC.

BZ#1244328

The iSCSI initiator name was the same for all Compute nodes in an

Overcloud, which causes live migration of instances to fail. This fix modifies the iSCSI initiator name during Overcloud deployment. Now live migration succeeds over iSCSI.

BZ#1278868

This enhancement adds support for the Nuage on highly available Overcloud environments. This includes Nuage-specific parameters in the director's Heat template collection, and environment files to enable the Nuage backend on Controller and Compute nodes.

BZ#1278879

This enhancement adds support for the Nuage metadata agent on the Overcloud. This includes parameters in the director's Heat template collection for the Nuage metadata agent.

BZ#1290050

The *ExtraConfig hiera data parameters did not work for non-Controller nodes. This is due to missing parameter definitions for non-Controller node types. This fix implements these parameters into the director's Heat template collection. Now the director writes the *ExtraConfig hiera data the the appropriate node types.

BZ#1290826

In Red Hat OpenStack Platform 7.0, Overclouds using a flat network had an additional Public IP created. This is no longer required but needs preservation for backwards compatibility. If upgrading from 7.0 and using a flat network topology, include the following environment file:

```
/usr/share/openstack-tripleo-heat-templates/environments/updates/update-from-publicvip-on-ctlplane.yaml
```

This environment file preserves the additional IP address from 7.0 and behaving as 'public_virtual_ip'.

BZ#1292562

High traffic Red Hat OpenStack Platform networks caused timeouts (specifically DNS timeouts) due to low maximum for netfilter connections tracking. This update increases the 'nf_conntrack_max' kernel parameter to 500000. This resolves the timeout issues.

BZ#1293473

This enhancement adds support to register Overcloud nodes to a Red Hat Satellite 5 server. Previous versions allowed registration only to a Red Hat Satellite 6 server. Now the director determines whether to register to a Red Hat Satellite 5 or Red Hat Satellite 6 server when using the '-reg-method satellite' option during Overcloud creation.

■

BZ#1295835

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#1296701

Swift caused deployment errors for an IPv6-based Overcloud due to problems with processing Swift's IPv6 addresses. This fix corrects how the IPv6 addresses are processed. Swift now deploys successfully.

BZ#1297850

Corosync failed to start in an IPv6-based Overcloud. This is due to a missing '--ipv6' option when the director tries to start Corosync. This fix adds this option to the Controller's Puppet manifest and also adds related parameters to the Heat template collection. Corosync now starts successfully in IPv6-based Overclouds.

BZ#1298197

This enhancement adds SSL support to the Overcloud's Public API. Users can now configure SSL on the Overcloud using the 'environments/enable-tls.yaml' from the director's Heat template collection. Copy and modify this environment file to suit your SSL requirements. For more information, see "6.2.7. Enabling SSL/TLS on the Overcloud" in the Director Installation and Usage guide for Red Hat OpenStack Platform 7.3.

BZ#1298198

The validation script for deployment testing node availability only supported IPv4. This caused connectivity checks for an IPv6-based Overcloud to fail. This fix modifies the validation script to detect whether the IP address is v4 or v6 and run the respective connectivity check commands. Now connectivity checks succeed for IPv6-based Overclouds.

BZ#1298222

The IPv6 network interface templates contained an error that set 'ExternalInterfaceDefaultRoute' to an IPv4 value. This fix corrects the error and sets the default route to an IPv6 value. The 'ExternalInterfaceDefaultRoute' now configures correctly.

BZ#1298506

Pacemaker failed to start in an IPv6-based Overcloud deployment due to using IPv4-based settings (/32) for the VIP netmask. This fix determines if the Overcloud uses IPv6 and sets the VIP netmasks to the appropriate values (/64 in most cases). Pacemaker now starts successfully in the Overcloud.

BZ#1298671

In an IPv6-based Overcloud, Galera failed to start due to issues with using an IPv6 address in configuration. This fix configures the 'bind-address' parameter to use the hostname, which all nodes should have in their '/etc/hosts' file. Galera now starts successfully in IPv6 Overclouds.

BZ#1299022

Non-Controller nodes reported package dependency issues due to delegating Puppet as the mechanism to update certain packages and excluding them from YUM updates. This fix sets all Non-Controller nodes to use Puppet as the update mechanism. Now packages on non-Controller nodes update without package dependency issues.

BZ#1299133

In an IPv6-based Overcloud, the director incorrectly parsed MariaDB DSN strings containing IPv6 addresses. This caused Puppet to report duplicate 'Mysql_database' resources due to all databases using the first bit grouping of the IPv6 address as the database name (e.g. 'fd00'). This fix adds logic to check if the string uses an IPv4 or IPv6 address and parse the string accordingly. Puppet no longer reports duplicate 'Mysql_database' resources.

BZ#1299265

'nova-consoleauth' failed to start due to how it parsed IPv6 addresses for 'memcached_servers' in 'nova.conf'. This fix corrects for the director's Heat template collection parses 'memcached_servers'. The 'nova-consoleauth' service now starts successfully.

BZ#1299294

In an IPv6-based Overcloud, RabbitMQ lacked some IPv6-specific options when starting. This caused RabbitMQ to fail on some nodes, which cause other services to fail due to Pacemaker constraints. This fix adds the IPv6-specific options. Now RabbitMQ and other services start successfully.

BZ#1299953

Overcloud deployments with IPv6 endpoints caused Glance to report a HTTP 500 error. This is due to how the director parsed IPv6 addresses. This

fix corrects how the director parses the IPv6 address for Glance. Glance now works in IPv6-based Overclouds.

BZ#1300678

The VNC servers on Compute nodes would bind to IPv4 addresses only. Users could not access VNC consoles in an IPv6 environment since the Internal API network used IPv6 endpoints. The fix allows the VNC servers to bind on IPv6 addresses when deploying an IPv6 overcloud. Now users can access instances through VNC consoles in an IPv6 deployment.

BZ#1300798

Setting a fixed IPv6 address for Overcloud networks failed due to Neutron not allowing fixed IP addresses in SLAAC mode. This fix changes the default IPv6 address mechanism to 'dhcpv6-stateful'. Now the director can configure the Overcloud using fixed IPv6 addresses.

BZ#1300800

The Puppet manifest in the Heat template collection started the HTTP service through systemd. This caused Pacemaker to fail starting the service. This fix modifies the manifest to only configure but not start the HTTP service so that Pacemaker can assume control over HTTP. Pacemaker now starts the service successfully.

BZ#1301015

In mixed environments where some networks use IPv4 addressing and others IPv6 addressing, the Overcloud used IPv6 CIDR for IPv4 VIPs too. The Overcloud deployment failed because Pacemaker refused to start the IPv4 VIPs. This fix adds functionality to identify the VIP type (IPv4 or IPv6) during deployment and adapt the appropriate CIDR. Each IPv4 and IPv6 VIP now uses the appropriate CIDR.

BZ#1301056

The Ceilometer Compute Agent in IPv6-based Overclouds could not reach the public endpoint and reported errors such as:

```
ConnectionError: ('Connection aborted.', gaierror(-2, 'Name or service not known'))
```

This fix switches the endpoint to 'internalURL' instead of 'publicURL'.

BZ#1301167

The 'router_delete_namespaces' (L3 agent) and 'dhcp_delete_namespaces' (DHCP agent) configuration settings defaulted to 'false' in the Red Hat OpenStack Platform 7 Puppet modules. This disabled cleanup for unused network namespaces, which was only required for older versions of Linux.

This fix sets the defaults for these parameters to 'true' in the Puppet modules so Red Hat OpenStack Platform takes advantage of network namespace cleanup.

BZ#1302593

The director inappropriately configured Ceph Storage IPv6, which caused deployment timeouts and Openstack deployment failure. This fix adds an input parameter (CephIPv6) in Heat templates, which sets the relevant Ceph config option (ms_bind_ipv6). Ceph Storage is now functional and the Openstack deployment completes when using IPv6 for the Storage network.

BZ#1303758

Compute nodes detected IPv6 router announcements (RA) on the Openstack's IPv6-based Tenant network. This caused the Compute nodes to receive IPv6 addresses and default routes on the 'qbr' interface. This fix sets the 'net.ipv6.conf.default.disable_ipv6' kernel parameter to 1 on all nodes. This disables automatic configuration from RAs and allows the director to define the IPv6 addresses and default route.

BZ#1304683

In an Openstack with HA Controller nodes, the 'cinder-volume' service might move to a new node. This causes problems modifying and deleting volumes due to a different hostname for the volume service. This fix sets a consistent hostname for the 'cinder-volume' service on all Controller nodes. Users can now modify and delete volumes on a HA Openstack without issue.

BZ#1304878

Heat used a 1MB payload size for returned output. YUM's output exceeded this limit if updating a high number of packages on the Openstack, which caused 'openstack overcloud update' to fail. This fix adds the '-q' option to YUM during an Openstack update. This option sets the output to quiet mode, which reduced the output. YUM's output no longer exceeds Heat's limit and 'openstack overcloud update' succeeds.

BZ#1305124

Compute nodes detected IPv6 router announcements (RA) on the Openstack's IPv6-based Tenant network. This caused problems with the IPv6 routing table, such as setting the default route to the Neutron router. This fix sets the 'net.ipv6.conf.default.accept_ra' kernel parameter to 0 on all nodes. Now the Compute node no longer accepts router announcements from the Tenant networks.

BZ#1305686

A bug in Heat caused validation of old parameters against the new

template. Upgrades of Overclouds from Red Hat OpenStack Platform 7.2 to Red Hat OpenStack Platform 7.3 failed with the error: "resources.SwiftDevicesAndProxyConfig: Property controller_swift_proxy_memcaches_v6 not assigned". This fix add defaults to Swift parameters in the director's Heat template collection, which resolves the error.

BZ#1306040

Ceilometer services failed to start in the Overcloud due to a incorrectly parsed IPv6 address from the director's Heat template collection. This fix correctly parses the IPv6 address. Ceilometer now starts correctly in an IPv6-based Overcloud.

BZ#1306623

When deploying an Overcloud with an external load balancer, the 'RedisVipPort' parameter resolved to the 'from_service.yaml' template in the director's Heat template collection. However, an issue with the 'ip_address_uri' output parameter in 'from_service.yaml' template provided the wrong value. This fix corrects the 'ip_address_uri' output parameter in 'from_service.yaml'. Now the 'from_service.yaml' template returns the correct value to 'RedisVipPort'.

6.7.5. os-cloud-config

BZ#1299604

Keystone endpoint creation failed due to incorrectly parsed IPv6 addresses. This fix modifies the Keystone client creation mechanism to correctly parse IPv6 addresses. 'os-cloud-config' now creates Keystone endpoints successfully.

BZ#1306022

Keystone client creation failed due to incorrectly parsed IPv6 addresses. This fix modifies the Keystone client creation mechanism to correctly parse IPv6 addresses. 'os-cloud-config' now creates Keystone clients successfully.

6.7.6. os-net-config

BZ#1298663

'os-net-config' only wrote IPv4 routes to /etc/sysconfig/network-scripts/route-*. However, IPv6 routes used /etc/sysconfig/network-scripts/route6-* files. This fix modifies 'os-net-config' to detect whether a route is IPv4 or IPv6 and write to the appropriate file. This allows 'os-net-config' to define IPv6 in the Linux configuration.

6.7.7. python-rdomanager-oscplugin

BZ#1296365

Multiple services attempted NTP configuration on the Overcloud and the last service configured it incorrectly. This caused time synchronization issues across all Overcloud nodes. As a workaround, delete `/usr/libexec/os-apply-config/templates/etc/ntp.conf` from all Overcloud nodes and re-run the deployment command to re-apply the puppet configuration. This is required for users updating from an older version of Red Hat OpenStack Platform to 7.3. This fix is not necessary on new 7.3 deployments. NTP now configures correctly.

6.7.8. rhel-osp-director

BZ#1296330

An issue with the OpenStack Platform director 7.2 ramdisk and kernel image caused provisioning failure with the following error:

```
mount: you must specify the filesystem type
Failed to mount root partition /dev/sda on /mnt/rootfs
```

This update reverts the ramdisk and kernel image to the OpenStack Platform director 7.1 images. Using these images, the director now provisions Overcloud nodes without failure.

NOTE: An alternative workaround is to disable localboot option for the different node types. For example, to disable localboot with Controller nodes, run:

```
$ nova flavor-key control unset capabilities:boot_option
```

BZ#1300264

'ceilometer-dbsync' fails in a highly available IPv6 Overcloud. This is due to how the director parsed IPv6 addresses for MongoDB. This fix corrects how the director parses IPv6 addresses. Now 'ceilometer-dbsync' runs successfully.

BZ#1300398

Horizon failed to load in IPv6 Overclouds due to issues with how the director detected and parsed IPv6 addresses for Memcached. This fix changes how the director's Heat template collection enables IPv6 addresses for Memcached. This includes a new parameter 'MemcachedIPv6' that defines if Memcached uses IPv4 or IPv6 addresses.

BZ#1301404

An SELinux issue stopped RabbitMQ from starting on IPv6-based Overclouds. This fix corrects the SELinux issue and RabbitMQ now starts successfully.

APPENDIX A. REVISION HISTORY

Revision 7.0.0-3	December 21, 2015	RHEL OpenStack Platform Docs Team
BZ#1293176 : Added VPNaaS and Time Series Database-as-a-Service to list of Technology Previews.		
Revision 7.0.0-2	December 18, 2015	RHEL OpenStack Platform Docs Team
Added descriptions for RHSA-2015:2650, RHBA-2015:2651, and RHBA-2015:2652.		
Revision 7.0.0-1	August 5, 2015	RHEL OpenStack Platform Docs Team
Adding link to certified plug-ins kbase.		
Revision 7.0.0-0	Mon Jun 9 2015	RHEL OpenStack Platform Docs Team
Initial revision for Red Hat Enterprise Linux OpenStack Platform 7.0		