Red Hat OpenStack Platform 11

Release Notes

Release details for Red Hat OpenStack Platform 11

Last Updated: 2019-04-24

Red Hat OpenStack Platform 11 Release Notes

Release details for Red Hat OpenStack Platform 11

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

Legal Notice

Copyright © 2016 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

CHAPTER 1. INTRODUCTION

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

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

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

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

1.1. ABOUT THIS RELEASE

This release of Red Hat OpenStack Platform is based on the OpenStack "Ocata" release. It includes additional features, known issues, and resolved issues specific to Red Hat OpenStack Platform.

Only changes specific to Red Hat OpenStack Platform are included in this document. The release notes for the OpenStack "Ocata" release itself are available at the following location: https://releases.openstack.org/ocata/index.html

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

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

To evaluate Red Hat OpenStack Platform, sign up at:

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

NOTE

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

1.2. REQUIREMENTS

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

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

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

> **NOTE**
>
> Prior to deploying Red Hat OpenStack Platform, it is important to consider the characteristics of the available deployment methods. For more information, refer to the Installing and Managing Red Hat OpenStack Platform.

1.3. DEPLOYMENT LIMITS

For a list of deployment limits for Red Hat OpenStack Platform, see Deployment Limits for Red Hat OpenStack Platform.

1.4. DATABASE SIZE MANAGEMENT

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

1.5. CERTIFIED DRIVERS AND PLUG-INS

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

1.6. CERTIFIED GUEST OPERATING SYSTEMS

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

1.7. BARE METAL PROVISIONING SUPPORTED OPERATING SYSTEMS

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

1.8. HYPERVISOR SUPPORT

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

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

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

1.9. CONTENT DELIVERY NETWORK (CDN) CHANNELS

This section describes the channel and repository settings required to deploy Red Hat OpenStack Platform 11.

You can install Red Hat OpenStack Platform 11 through the Content Delivery Network (CDN). To do so, configure **subscription-manager** to use the correct channels.

WARNING

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

Run the following command to enable a CDN channel:

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

Run the following command to disable a CDN channel:

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

Table 1.1. Required Channels

Channel	Repository Name
Red Hat Enterprise Linux 7 Server (RPMS)	**rhel-7-server-rpms**
Red Hat Enterprise Linux 7 Server - RH Common (RPMs)	**rhel-7-server-rh-common-rpms**
Red Hat Enterprise Linux High Availability (for RHEL 7 Server)	**rhel-ha-for-rhel-7-server-rpms**
Red Hat OpenStack Platform 11 for RHEL 7 (RPMs)	**rhel-7-server-openstack-11-rpms**
Red Hat Enterprise Linux 7 Server - Extras (RPMs)	**rhel-7-server-extras-rpms**

Table 1.2. Optional Channels

Channel	Repository Name
Red Hat Enterprise Linux 7 Server - Optional	`rhel-7-server-optional-rpms`
Red Hat OpenStack Platform 11 Operational Tools for RHEL 7 (RPMs)	`rhel-7-server-openstack-11-optools-rpms`

Channels to Disable

The following table outlines the channels you must disable to ensure Red Hat OpenStack Platform 11 functions correctly.

Table 1.3. Channels to Disable

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

WARNING

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

1.10. PRODUCT SUPPORT

Available resources include:

Customer Portal

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

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

- Support case management.

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

Mailing Lists

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

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

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

CHAPTER 2. TOP NEW FEATURES

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

2.1. RED HAT OPENSTACK PLATFORM DIRECTOR

This section outlines the top new features for the director.

Composable Services Upgrades

Each composable service template now contains logic to upgrade the service across major releases. This provides a mechanism to accommodate upgrades through the custom role and composable service architecture.

Deployment on Pre-Provisioned Infrastructure

The director now configures Red Hat OpenStack Platform on existing systems running the latest release of Red Hat Enterprise Linux (for the initial release of Red Hat OpenStack Platform 11, this is version 7.3). This means you can provision systems outside the standard director tools but still use the director to configure these systems to run Red Hat OpenStack Platform while using the custom role and composable service architecture.

Support for Standalone Ironic Role

The director can now deploy OpenStack Bare Metal Provisioning (ironic) as a standalone role on the overcloud. In previous version, the **`Ironic`** role required inclusion on a Split Systemd Controller. Now you can deploy the role on its own.

Dynamic Ansible Inventory

The director now provides the **`tripleo-ansible-inventory`** command that generates an inventory of hosts in the environment. Use this for running Ansible automation tasks on groups of hosts.

2.2. BLOCK STORAGE

NFS Snapshots

The NFS back end driver for the Block Storage service now supports snapshots.

2.3. COMPUTE

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

Placement API Service

This release includes the placement API service. This service is a separate REST API stack and data model that tracks the inventory and usage of resource providers (Compute nodes).

You must deploy the placement API service after upgrading to the Red Hat OpenStack Platform 10 release but before upgrading to the Red Hat OpenStack Platform 11 release. This is so the **`nova-compute`** service's resource tracker can populate the resource provider inventory and allocation information that the **`nova-scheduler`** service uses in Red Hat OpenStack Platform 11.

VLAN Metadata Exposure

The SR-IOV physical function is now exposed to the guests by providing VLAN tags in the metadata. This feature expands the role tagging functionality of the devices that was introduced in previous releases.

The following example shows the possible metadata structure to illustrate how VLAN tags are passed through.

```
{"devices": [{
  "type": "nic",
  "bus": "pci",
  "address": "0000:00:02.0",
  "mac": "01:22:22:42:22:21",
  "tags": ["nfvfunc1"]
  "vlans":[300,1000]
  }]
}
```

EC2 API Deployment and Configuration

OpenStack Compute now provides EC2 API support as a standalone service that consumes **nova**. The Red Hat OpenStack Platform director can now deploy this service.

2.4. DASHBOARD

This section outlines the top new features for the Dashboard.

Improved Parity with Core OpenStack Services

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

Improved User Experience

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

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

2.5. IDENTITY

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

Documentation - Keystone Federation with RH-SSO

Detailed documentation for director-based deployments of Identity Service (keystone) backed by Red Hat Single Sign On. This guide describes SAML-based federation and uses Red Hat Single Sign-On (RH-SSO) as the external identity provider: Federate with Identity Service

Domain-Specific Roles

Allows role definition to be limited to a specific domain, or a project with a domain. Domain-specific roles grant you more granular control when defining rules for roles, allowing the roles to act as aliases for the existing `prior` roles.

Implied Roles

Implied roles means that your role assignments are processed cumulatively. For example, if a user has the **`admin`** role on a project, they would also be a **`_member_`** of that project, even though the **`_member_`** role was not explicitly assigned. This is because an inference rule can be set saying that assignment of one role implies the assignment of another. This feature is expected to make role management much easier for admins.

2.6. IMAGE SERVICE

Improved Image Signing and Trust

The Image service features improved handling of authentication tokens, thereby ensuring that image uploads from trusted users are handled correctly. In previous releases, it was possible for a user's authentication token to expire while uploading a large image, thereby causing the upload to fail.

2.7. OPENSTACK NETWORKING

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

VLAN-Aware VMs

Instances can now send and receive VLAN-tagged traffic over a single vNIC. This ability is particularly useful for NFV applications (VNFs) that expect 802.1q VLAN-tagged traffic, allowing multiple customers/services to be served by a single vNIC. This implementation has full support with OVS-based and OVS-DPDK-based networks.

WARNING

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

2.8. SHARED FILE SYSTEM

Enhanced User Interface

This release features several improvements to the dashboard interface of the Shared File System service. These improvements include an improved Shared drop-down menu for selecting back ends.

In addition, you can now disable the creation of Public Shares by manually editing the local settings of your dashboard service.

2.9. TELEMETRY

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

Ceilometer

To provide better performance and scalability, the ceilometer API has been replaced by *gnocchi* and *aodh*, and is now considered *deprecated*. The ceilometer events API and its code have been moved to a new component called *panko*.

The nova instance discovery, that was based on polling nova API, was very resource consuming. Newly, it has been optimized to rely on libvirt, which significantly improves the performance.

Gnocchi

Gnocchi provides a new *collectd* plugin that stores metrics generated by collectd.

Panko

Panko is a new component that replaces the ceilometer events and its API.

2.10. HIGH AVAILABILITY

This section outlines the top new features for high availability.

Composable High Availability Services

The Red Hat OpenStack Platform director now opens the composable service architecture to include high availability services. This means users can split high availability services from the Controller node or scale services with dedicated custom roles. This includes the following high availability services:

- Load Balancer (HAProxy)
- Database (MariaDB/Galera)
- Messaging (RabbitMQ)
- Redis
- Block Storage (cinder) Volume
- Block Storage (cinder) Backup
- OpenStack Shared File Systems (manila)

2.11. OPERATIONS TOOLING

This section outlines the top new features for operations tooling.

Red Hat OpenStack Platform 11 includes full support for performance monitoring (collectd), log aggregation (fluentd), and availability monitoring (sensu). These agents are called *composable services* by Red Hat OpenStack Platform director, and are configured with Heat templates during installation.

Performance Monitoring

There is now full support for *collectd* clients to monitor performance in the overcloud.

Common Logging

fluentd collects log data and then forwards the logs from the overcloud nodes to a remote fluentd instance.

Availability Monitoring

The *sensu* agent runs scripts to check whether certain conditions are met, then delivers the results to the server.

2.12. BARE METAL PROVISIONING SERVICE

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

Graceful Shutdown and NMI

This release provides support for graceful shutdown and nonmaskable interrupts (NMI) for bare metal nodes. Graceful shutdown provides a safe way to power off the bare metal node using the API, which is useful when SSH connectivity is unavailable. NMI assists in accessing the bare metal node for troubleshooting and core dumps.

LLDP Data Extraction

The Bare Metal Provisioning service can now extract LLDP data, including information about the attached switch port, during the inspection of overcloud nodes. The information is extracted by querying the Swift object associated with each node.

VirtualBMC and IPMI

This release introduces the VirtualBMC proxy tool to control virtual machine power for bare metal nodes using the IPMI protocol. You can use this feature with the pxe_ipmitool driver to replace the deprecated pxe_ssh drivers to test bare metal deployments in a virtual environment. VirtualBMC and the pxe_ipmitool allow you to use the same drivers for test and deployment of bare metal nodes.

2.13. OPENSTACK INTEGRATION TEST SUITE SERVICE

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

Identity Service Clients as Libraries

New service clients have been added to the library interface, so that other projects can use these modules. This includes stable modules such as identity, groups, trusts and users.

Volume Service Clients as Libraries

Volume service clients have been added to the library interface, so that other projects can use these libraries. This includes clients for backups, encryption, QoS and snapshots.

2.14. OPENSTACK DATA PROCESSING SERVICE

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

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

This release adds support for MapR 5.2 and 5.1 plugins.

2.15. TECHNOLOGY PREVIEWS

This section outlines features that are in technology preview in Red Hat OpenStack Platform 11.

NOTE

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

2.15.1. New Technology Previews

The following new features are provided as technology previews:

Benchmarking Service - Introduction of a new plug-in type: Hooks

Allows test scenarios to run as iterations, and provides timestamps (and other information) about executed actions in the rally report.

Benchmarking Service - New Scenarios

Benchmarking scenarios have been added for nova, cinder, magnum, ceilometer, manila, and neutron.

Benchmarking Service - Refactor of the Verification Component

Rally Verify is used to launch Tempest. It was refactored to cover a new model: verifier type, verifier, and verification results.

Block Storage - Highly Available Active-Active Volume Service

In previous releases, the openstack-cinder-volume service could only run in Active-Passive HA mode. Active-Active configuration is now available as a technology preview with this release. This configuration aims to provide a higher operational SLA and throughput.

IMPORTANT

The active-active volume functionality is available only if you already have a Block Storage driver that supports active-active configuration. This driver is not provided as a part of this release.

Block Storage - RBD Cinder Volume Replication

The Ceph volume driver now includes RBD replication, which provides replication capabilities at the cluster level. This feature allows you to set a secondary Ceph cluster as a replication device; replicated volumes are then mirrored to this device. During failover, all replicated volumes are set to 'primary', and all new requests for those volumes will be redirected to the replication device.

To enable this feature, use the parameter replication_device to specify a cluster that the Ceph back end should mirror to. This feature requires both primary and secondary Ceph clusters to have RBD mirroring set up between them. For more information, see http://docs.ceph.com/docs/master/rbd/rbd-mirroring/.

At present, RBD replication does not feature a failback mechanism. In addition, the freeze option does not work as described, and replicated volumes are not automatically attached/detached to the same instance during failover.

CephFS Integration - CephFS Native Driver Enhancements

The CephFS driver is still available as a Technology Preview, and features the following enhancements:

- Read-only shares
- Access rules sync
- Backwards compatibility for earlier versions of **`CephFSVolumeClient`**

Link Aggregation for Bare Metal Nodes

This release introduces link aggregation for bare metal nodes. Link aggregation allows you to configure bonding on your bare metal node NICs to support failover and load balancing. This feature requires specific hardware switch vendor support that can be configured from a dedicated neutron plug-in. Verify that your hardware vendor switch supports the correct neutron plug-in.

Alternatively, you can manually preconfigure switches to have bonds set up for the bare metal nodes. To enable nodes to boot off one of the bond interfaces, the switches need to support both LACP and LACP fallback (bond links fall back to individual links if a bond is not formed). Otherwise, the nodes will also need a separate provisioning and cleaning network.

2.15.2. Previously Released Technology Previews

The following features remain as technology previews:

Benchmarking Service

Rally is a benchmarking tool that automates and unifies multi-node OpenStack deployment, cloud verification, benchmarking and profiling. It can be used as a basic tool for an OpenStack CI/CD system that would continuously improve its SLA, performance and stability. It consists of the following core components:

1. Server Providers - provide a unified interface for interaction with different virtualization technologies (LXS, Virsh etc.) and cloud suppliers. It does so via ssh access and in one L3 network
2. Deploy Engines - deploy an OpenStack distribution before any benchmarking procedures take place, using servers retrieved from Server Providers
3. Verification - runs specific set of tests against the deployed cloud to check that it works correctly, collects results & presents them in human readable form
4. Benchmark Engine - allows to write parameterized benchmark scenarios & run them against the cloud.

Cells

OpenStack Compute includes the concept of Cells, provided by the nova-cells package, for dividing computing resources. In this release, Cells v1 has been replaced by Cells v2. Red Hat OpenStack Platform deploys a "cell of one" as a default configuration, but does not support multi-cell deployments at this time.

CephFS Native Driver for Manila

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

Containerized Compute Nodes

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

DNS-as-a-Service (DNSaaS)

Red Hat OpenStack Platform 11 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. DNSaaS includes integration with the Bind9 back end.

Firewall-as-a-Service (FWaaS)

The Firewall-as-a-Service plug-in adds perimeter firewall management to OpenStack Networking (neutron). FWaaS uses iptables to apply firewall policy to all virtual routers within a project, and supports one firewall policy and logical firewall instance per project. FWaaS operates at the perimeter by filtering traffic at the OpenStack Networking (neutron) router. This distinguishes it from security groups, which operate at the instance level.

Google Cloud Storage Backup Driver (Block Storage)

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

Object Storage Service - At-Rest Encryption

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

Object Storage Service - Erasure Coding (EC)

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

OpenDaylight Integration

Red Hat OpenStack Platform 11 includes a technology preview of integration with the OpenDaylight SDN controller. OpenDaylight is a flexible, modular, and open SDN platform that supports many different applications. The OpenDaylight distribution included with Red Hat OpenStack Platform 11 is limited to the modules required to support OpenStack deployments using NetVirt, and is based on the upstream Boron version.

For more information, see the Red Hat OpenDaylight Product Guide and the OpenDaylight and Red Hat OpenStack Installation and Configuration Guide.

Open vSwitch Firewall Driver

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

Real Time KVM Integration

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

Red Hat SSO

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

VPN-as-a-Service (VPNaaS)

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

IMPORTANT

VPNaaS is deprecated in Red Hat OpenStack Platform 11 and is planned to be removed in Red Hat OpenStack Platform 12.

CHAPTER 3. RELEASE INFORMATION

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

Notes for updates released during the support lifecycle of this Red Hat OpenStack Platform release will appear in the advisory text associated with each update.

3.1. RED HAT OPENSTACK PLATFORM 11 GA

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

3.1.1. Enhancements

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

BZ#962864

> This update adds a new user details page to the dashboard. As a result you can click on the user ID in the Action Log to be taken directly to the user details page.

BZ#962864

> This update adds a new user details page to the dashboard. As a result you can click on the user ID in the Action Log to be taken directly to the user details page.

BZ#1197163

> The Time Series Database as a Service (gnocchi) and Aodh API endpoints now expose a `/healthcheck` HTTP endpoint on the REST API. Requesting this endpoint allows you to check the status of the service, and does not require authentication.

BZ#1242422

> Automatic fencing setup can be used in director for easier High Availability deployments and upgrades. To benefit from the new feature, use the 'overcloud generate fencing' command.

BZ#1271019

> The administrator needs to record the user credentials during the volume transfer operation and doing so by hand is inconvenient.
>
> With this update, a new button to download the credentials has been

> added to the volume transfer screen for saving the information easily. This allows the administrators to download and save a CSV file locally on their computer with the click of a button.

BZ#1325861

> This enhancement adds the ability to automatically reschedule load balancers from LBaaS agents that the server detects as dead. Previously, load balancers could be scheduled and realized across multiple LBaaS agents, however if a hypervisor died, the load balancers scheduled to that node would cease operation. With this update, these load balancers will be automatically rescheduled to a different agent. This feature is disabled by default and managed using `allow_automatic_lbaas_agent_failover`.

BZ#1326224

> This enahncement implements the 'ProcessMonitor' class in the 'HaproxyNSDriver' class (v2), This class utilizes the 'external_process' module to monitor and respawn HAProxy processes if and when needed. The LBaaS agent (v2) loads 'external_process' related options and take a configured action when HAProxy dies unexpectedly.

BZ#1337664

> This update adds support for the version 5.1.0 MapR plugin.

BZ#1377867

> A disk can be in a variety of states which may cause director to fail when attempting to make the disk a Ceph OSD. In previous releases, a user could run a first-boot script to erase the disk and set a GPT label required by Ceph. With this release, a new default setting in Ironic will erase the disks when a node is set to available and a change in puppet-ceph will give the disk a GPT label if there is no GPT label on the disk.

BZ#1386249

> This update provides enhancements to the CephFS Native Driver in conjunction with the core OpenStack File Share Service (manila) infrastructure. The CephFS Native driver now supports read-only shares and includes improvements to access rule update recovery mode.

BZ#1388171

> To avoid memory bloat issues in the nova-api workers, pagination logic has been added to the simple-tenant-usage API extension.

BZ#1393893

> With this enhancement, you can now enable the creation of non-public

```
shares in the Dashboard.
You can configure Dashboard to hide the checkbox that enables users to
mark shares as public during the creation process. The default option
will be to create shares as private without the box checked.
```

BZ#1396794

```
With this enhancement, `glance-manage db purge` can now remove rows that
are less than one day old. This was added because operators may need to
run this operation on a regular basis.
As a result, the value of the `age_in_days` option can be set to `0`.
```

BZ#1413980

```
This release features the necessary puppet modules for deploying CephFS.
This allows you to deploy the OpenStack Shared File System service
(openstack-manila) with a CephFS back-end through the director.
```

BZ#1421554

```
This enhancement makes Time Series Database as a Service (gnocchi)
available in the undercloud. Gnocchi provides metrics backend for
Telemetry in OpenStack, and is enabled by default with the
`enable_telemetry` flag set to `true`. All telemetry services can be
disabled by setting `enable_telemetry=false` in `undercloud.conf`.
```

3.1.2. Release Notes

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

BZ#1352922

```
This release adds pagination support to avoid resource-consuming usage
requests on systems with a large number of instances. The v2.40
microversion of the nova API simple-tenant-usage endpoints use new
optional query parameters 'limit' and 'marker' for pagination. The
'marker' option sets the starting point and the 'limit' option sets the
number of records to be displayed after the starting point. If 'limit'
is not set, nova will use the configurable 'max_limit' (1000 by
default). Although older microversions will not accept these new query
parameters, they will start to enforce the max_limit and results may be
truncated as a result. Consider using the new microversion to avoid DoS-
like usage requests and potentially truncated responses.
```

BZ#1383199

```
With this update, different domain names can be used for the public and
the internal networks.  To set domain names per network, use the
following heat template parameters:
```

```
* CloudName: The DNS name of this cloud, for example:
      'ci-overcloud.tripleo.org'
* CloudNameInternal: The DNS name of this cloud's internal API endpoint,
for example:
      'ci-overcloud.internalapi.tripleo.org'
* CloudNameStorage: The DNS name of this cloud's storage endpoint, for
example:
      'ci-overcloud.storage.tripleo.org'
* CloudNameStorageManagement: The DNS name of this cloud's storage
management endpoint, for example:
      'ci-overcloud.storagemgmt.tripleo.org'
* CloudNameCtlplane: The DNS name of this cloud's control plane
endpoint, for example:
      'ci-overcloud.management.tripleo.org'
```

BZ#1386309

```
With this update, the user interface is now partially internationalized
and available in Japanese and Simplified Chinese. Note that only the
interface itself is internationalized at this stage. The strings that
come from other services such as the parameters, templates, and
environments for the TripleO Heat Templates, the validations, and the
notifications are not yet internationalized.
```

BZ#1399816

```
Recent enhancements to the director requires changes to network
interface configuration templates. The NIC configuration templates now
use a script that calls the 'os-net-config utility' to configure the
network on the overcloud nodes. There are three major changes to the NIC
config templates:

* The 'OsNetConfigImpl' resource changed from a
'OS::Heat::StructuredConfig' resource type to
'OS::Heat::SoftwareConfig'. In addition, the resource now stores the
'network_config' property as a blob of text and passes the blob to the
'run-os-net-config.sh' script using the 'str_replace' (string replace)
function. For example:

----
resources:
  OsNetConfigImpl:
    type: OS::Heat::SoftwareConfig
    properties:
      group: script
      config:
        str_replace:
          template:
            get_file: ../../scripts/run-os-net-config.sh
          params:
            $network_config:
              network_config:
----

* The {get_input: <input>} constructor defined a default external bridge
```

```
and interface. Now there are two special string values that substitute
for the external bridge and interface. These are 'bridge_name' and
'interface_name' respectively. Instead of using '{get_input:
bridge_name}' or '{get_input: interface_name}', use 'bridge_name' or
'interface_name'. For example:

----
              - type: ovs_bridge
                name: {get_input: bridge_name}
----
becomes:
----
              - type: ovs_bridge
                name: bridge_name
----

* The 'network_config' no longer uses curly braces. Instead, the
{get_param: <param>} construct moves to a sub-level underneath the value
being defined. For example:
----
                dns_servers: {get_param: DnsServers}
----
becomes:
----
                dns_servers:
                  get_param: DnsServers
----

See more examples in the "Network Isolation" chapter of the Advanced
Overcloud Customizations guide.
```

BZ#1427507

```
With this update, Wake-On-LAN and AMT drivers have been removed from
Ironic as they do not have, and are not planned to have, a third-party
CI. They are still available from the unsupported ironic driver
collection, found in the ironic-staging-drivers repository. If your
ironic installation is using any driver based on those, you must install
ironic-staging-drivers and change the driver on the affected nodes
according to following list:

agent_amt -> pxe_amt_agent
pxe_amt -> pxe_amt_iscsi
agent_wol -> pxe_wol_agent
pxe_wol -> pxe_wol_iscsi
```

BZ#1431556

```
Because SELinux policies concerning launching instances with DPDK
enabled are incomplete, launching instances using DPDK with SELinux in
enforcing mode will cause the launch to fail and AVC denials will appear
in /var/log/audit/audit.log* concerning openvswitch and svirt.

As a workaround, set SELinux to permissive on each compute node where
DPDK is utilized as documented in section 4.4.1.2 here:
```

```
https://access.redhat.com/documentation/en-
US/Red_Hat_Enterprise_Linux/7/html/SELinux_Users_and_Administrators_Guid
e/sect-Security-Enhanced_Linux-Working_with_SELinux-
Changing_SELinux_Modes.html#sect-Security-Enhanced_Linux-
Enabling_and_Disabling_SELinux-Disabling_SELinux

This will allow DPDK-enabled virtual machines to launch. This is a
workaround and is expected to be temporary while the issue is
investigated further.
```

BZ#1451714

```
Problem in detail:
In OSP10 (OvS2.5), following are the issues:
1) tuned is configured with wrong set of CPUs. Expected configuration is
NeutronDpdkCoreList + NovaVcpuPinSet, but it has been configured as
HostCpusList.
2) In post-config, the -l of DPDK_OPTIONS is set as 0 and
NeutronDpdkdCoreList is configured as pmd-cpu-mask

What needs to be corrected after update, manually?
1) Add the list of cpus to be isolated, which is NeutronDpdkCoreList +
NovaVcpuPinSet to the tuned conf file.

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

2) lcore mask after the update will be set to 0. Get the cpu mask with
get_mask code from the first-boot script [1].
LCORE_MASK="<mask value output of get_mask"
ovs-vsctl --no-wait set Open_vSwitch . other-config:dpdk-lcore-
mask=$LCORE_MASK
```

3.1.3. Known Issues

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

BZ#1225069

```
For security reasons, the Overcloud only allows SSH key-based access by
default. You can set a root password on the disk image for the overcloud
using the virt-customize tool, which is found in the Red Hat Enterprise
Linux Extras channel. After installing the tool and downloading the
Overcloud images, use the following command to change the root password:

$ virt-customize -a overcloud-full.qcow2 --root-password password:
<my_root_password>

Perform this operation prior to uploading the images into glance with
the "openstack overcloud image upload" command.
```

BZ#1227955

> Only NICs which have a live connection to a switch port will be counted in the numbered NIC abstractions (nic1, nic2, etc). As a workaround, the director includes a script that pings the first controller on all interfaces from each node. If a node has a link disconnected at deployment time, this is detected so that it can be corrected. Another possible workaround is to use a mapping file for each host, which has a mapping of NIC number to physical NIC. If one or more Overcloud nodes do not configure correctly due to a down link, this is now detected and the Overcloud can redeploy.

BZ#1243109

> Discovery fails if multiple network interfaces on a node are connected to the Provisioning network. Only one interface can connect to the Provisioning network. This interface cannot be part of a bond.

BZ#1247019

> Pacemaker continuously crashes when the fencing device name and the host name are the same. To avoid this problem, add the "fence-" prefix or the "-fence" suffix to the name of the fencing device. With the names configured like this, the cluster works without errors.

BZ#1369591

> When you enter the `openstack stack delete` command while the nodes are still being deployed, it is possible that they remain in that being-deployed state and will not be deleted. Consequently, you will not be able to deploy any new nodes onto the space, blocked by the undeleted nodes. To avoid this situation, wait until all the nodes have been deployed completely, before entering the `openstack stack delete` command. Alternatively, you can manually delete the nodes with the `nova delete <node>` command.

BZ#1384126

> The 'openstack overcloud stack update' command has a lengthy startup time. The command can take a couple of minutes for an indication the command is working. This is normal behavior.

BZ#1385338

> To implement the security groups trunk feature with neutron-openvswitch-agent, openvswitch firewall driver is required. This driver currently contains a bug 1444368 where ingress traffic is wrongly matched if there are two ports with same MAC address on different network segment on the same compute node.
>
> As a result, if a subport has the same MAC address as its parent port, ingress traffic won't be matched correctly for one of the ports.

> A workaround to achieve correctly handled traffic is to disable port-security on the parent port and subports.
>
> For example, to disable port security on port with UUID 12345, you need to remove security groups associated with the port:
>
> ```
>  openstack port set --no-security-group --disable-port-security 12345
> ```
>
> Note that no security groups rules will be applied to that port and traffic will not be filtered or protected against ip/mac/arp spoofing.

BZ#1392155

> A race condition may occur between Puppet and MongoDB services. Consequently, scaling out nodes running the MongoDB database fails and the overcloud stack will not update. Running the same deployment command again makes the MongoDB nodes scale out successfully.

BZ#1409097

> Currently, the Red Hat OpenStack Platform director 10 with SRIOV overcloud deployment fails when using the NIC IDs (for example, nic1, nic2, nic3 and so on) in the compute.yaml file.
>
> As a workaround, you need to use NIC names (for example, ens1f0, ens1f1, ens2f0, and so on) instead of the NIC IDs to ensure the overcloud deployment completes successfully.

BZ#1430757

> Under certain circumstances, such as when OVS is upgraded, or when the network service is restarted, the OVS bridges were torn down and rebuilt. Consequently, the existing network flows were interrupted when this happened, causing network traffic to stop forwarding until the flows were rebuilt. This can take some time in a complex deployment.
>
> In order to avoid any possible downtime, the control plane networks should not be placed on an OVS bridge. The Control Plane (Provisioning), Internal API, and Storage Management networks should instead be dedicated interfaces or VLAN interfaces that are not on a bridge. For instance, one interface or bond could contain the control plane VLANs, while another interface or bond can be placed on an OVS bridge for tenant network data.
>
> As long as the control plane interfaces are not on an OVS bridge, any network downtime will be limited to the Tenant data plane.

BZ#1437566

> Processing parameters, environments and templates is slightly different on the CLI than it is in the UI.
> Consequently, the passwords generated automatically by the UI cannot be changed from the templates. If you want to use custom passwords, you have to set them manually in the UI as a parameter in the overall deployment configuration or by editing the role card.

> Alternatively, you can create a plan without auto-generated passwords by entering the '$ openstack overcloud plan create <my_plan> --disable-password-generation' on the CLI. You will have to provide the passwords explicitly by using templates or manually through the UI.

BZ#1440273

> Running the 'openstack overcloud deploy' command replaces the default 'overcloud' plan. If a user creates an overcloud with the CLI, deletes it, then creates a new overcloud with the web UI, the web UI uses the 'overcloud' plan from the CLI deployment. This can cause the web UI to include unwanted parameters from the previous overcloud deployment. As a workaround:
>
> 1. Ensure the 'user-environment.yaml' environment file is disabled when deploying a new overcloud.
> 2. Upload a new version of the plan (from the '/usr/share/openstack-tripleo-heat-template').

BZ#1440276

> The director requires nodes in a managed state before running introspection. Although newly registered nodes are set to 'manageable' in the director's web UI, no option currently exists in the web UI to switch the nodes back to 'manageable' in case users require introspection at a later date. As a workaround, use the 'openstack baremetal node manage' command to switch the nodes to 'manageable' state.

BZ#1441393

> Invalid cache files may cause os-collect-config to report 'ValueError: No JSON object could be decoded' and the service will fail to start. The cache files located in '/var/lib/os-collect-config/' should be valid json files. If they are are of size 0 or contain invalid json, remove the invalid files from '/var/lib/os-collect-config', otherwise they may prevent os-collect-config from starting.

BZ#1445886

> Customers who upgraded from Red Hat OpenStack Platform 9 to version 10 are advised to wait for the first asynchronous release before upgrading to Red Hat OpenStack Platform 11, which will fix this known issue. The first asynchronous release typically happens within a few days of the GA release.

BZ#1445905

> In Highly Available IPv6 deployments, virtual IPs used for RabbitMQ may move between controller hosts during an upgrade. A bug in the creation of these IPv6 IPs causes them to be used as source addresses for RabbitMQ's connections. As a result, RabbitMQ will crash and may be unable to automatically recover its cluster.

> To return to normal operation, restart RabbitMQ on the affected controller hosts, as well as any services which depend on RabbitMQ and do not automatically reconnect.

BZ#1445917

> Customers who upgraded from Red Hat OpenStack Platform 9 to version 10 are advised to wait for the first asynchronous release before upgrading to Red Hat OpenStack Platform 11, which will fix this known issue. The first asynchronous release typically happens within a few days of the GA release.

BZ#1446825

> A design flaw issue was found in the Red Hat OpenStack Platform director use of TripleO to enable libvirtd based live migration. TripleO did not have support for secure live migration and no additional steps were taken to lock-down the libvirtd deployment by director. Libvirtd is deployed by default (by director) listening on 0.0.0.0 (all interfaces) with no-authentication or encryption. Anyone able to make a TCP connection to any compute host IP address, including 127.0.0.1, other loopback interface addresses or in some cases possibly addresses that have been exposed beyond the management interface, could use this to open a virsh session to the libvirtd instance and gain control of virtual machine instances or possibly take over the host.
>
> Note that without the presence of additional flaws, this should not be accessible from tenant or external networks.
>
> Users who are upgrading to Red Hat OpenStack Platform 11 from Red Hat OpenStack Platform 10 should first apply the relevant update that resolves this issue.
>
> Red Hat OpenStack Platform 11 already contains this update as of general availability and no subsequent update is required.
>
> For more information about this flaw and the accompanying resolution, see https://access.redhat.com/solutions/3022771.

BZ#1447731

> If using a standalone Keystone node on OpenStack Platform 10, 'openstack-gnocchi-statsd' does not start correctly. This is because 'gnocchi' and 'keystone' services activate on the same step, which causes a race condition. 'gnocchi' fails authentication but does not retry. This issue is addressed in BZ#1447422.
>
> The failed 'openstack-gnocchi-statsd' service causes OpenStack Platform 11 pre-upgrade check to fail, which means upgrades from OpenStack Platform 10 to 11 also fail if using a standalone Keystone role. As a work around, restart 'openstack-gnocchi-statsd' service on the overcloud before starting any upgrade steps. This enables the service correctly

> and allows for a successful upgrade. This only applies to OpenStack Platform 10 with a standalone Keystone node and does not impact other upgrade scenarios.

BZ#1463058

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

BZ#1321179

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

3.1.4. Deprecated Functionality

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

BZ#1256912

> The image_path parameter is no longer used. This update removes it from the undercloud configuration file.

BZ#1426917

> The VPNaaS feature is deprecated with Red Hat OpenStack Platform 11 and is expected to be removed with Red Hat OpenStack Platform 12.

BZ#1426919

> Neutron's Linux Bridge ML2 driver and agent are being deprecated with Red Hat OpenStack Platform 11 and are expected to be removed with Red Hat OpenStack Platform 12. The Open vSwitch (OVS) plug-in is the one deployed by default by the OpenStack Platform director, and is recommended by Red Hat for general usage.

BZ#1432458

> The Ceilometer API service is deprecated in Red Hat OpenStack Platform 11. It is replaced by the Gnocchi, Aodh, and Panko APIs respectively. Users should begin to move away from the Ceilometer API and use the service APIs instead. In Red Hat OpenStack Platform 11, the Ceilometer API is installed and configured by default. In future releases, this API will be disabled by default, with the option to enable it only if required.

BZ#1461990

> As of this release, Glance API V1 is no longer supported or available.

3.2. RED HAT OPENSTACK PLATFORM 11 MAINTENANCE RELEASES

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

3.2.1. Enhancements

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

BZ#962864

> This update adds a new user details page to the dashboard. As a result you can click on the user ID in the Action Log to be taken directly to the user details page.

BZ#962864

> This update adds a new user details page to the dashboard. As a result you can click on the user ID in the Action Log to be taken directly to the user details page.

BZ#1378993

> This enhancement provides configuration for OpenStack Load Balancing as a Service (octavia) through the Red Hat OpenStack Platform director.

BZ#1439855

> A high memory consumption was observed, especially in large environments, which often led to out-of-memory issues. The main culprit was neutron-ns-metadata-proxy process, responsible for proxying metadata requests from the VM to Nova.
>
> neutron-ns-metadata-proxy is now replaced by haproxy which has a more lightweight memory footprint.

BZ#1498108

> The OS::Nova::ServerGroup resource now allows you to use the 'soft-affinity' and 'soft-anti-affinity' policies. This is in addition to the 'affinity' and 'anti-affinity' policies.

3.2.2. Release Notes

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

BZ#1438828

To use SR-IOV physical function (PF) and virtual functions (VFs) in the same environment, add the 'nm_controlled' and 'hotplug' parameters to the SR-IOV PF configuration in your compute.yaml heat template:

```
 -type: interface
  name: nic6
  use_dhcp: false
  nm_controlled: true
  hotplug: true
```

When an OpenStack instance that was using a direct physical function is destroyed, the PCI device is released back to OpenStack and the host system. The root PCI device is then configured to support the number of virtual functions configured during deployment. This process involves the coordination of the host operating system, NetworkManager and OpenStack and may require a short interval of time before the virtual functions are available for use.

BZ#1441811

Red Hat OpenStack Platform 11 ships with many fewer API workers for each service than in the previous version. Having fewer workers lowers the noise, but also lowers the performance of API response times.

3.2.3. Known Issues

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

BZ#1441393

Invalid cache files may cause os-collect-config to report 'ValueError: No JSON object could be decoded' and the service will fail to start. The cache files located in '/var/lib/os-collect-config/' should be valid json files. If they are are of size 0 or contain invalid json, remove the invalid files from '/var/lib/os-collect-config', otherwise they may prevent os-collect-config from starting.

BZ#1445905

In Highly Available IPv6 deployments, virtual IPs used for RabbitMQ may move between controller hosts during an upgrade. A bug in the creation of these IPv6 IPs causes them to be used as source addresses for RabbitMQ's connections. As a result, RabbitMQ will crash and may be unable to automatically recover its cluster.

To return to normal operation, restart RabbitMQ on the affected controller hosts, as well as any services which depend on RabbitMQ and do not automatically reconnect.

BZ#1445917

> A database upgrade issue hindered the OpenStack Networking (neutron) upgrade. This hindered the upgrade from OpenStack Platform 10 to 11. This fix corrects the neutron database upgrade issue.
>
> This issue affects customers who previously upgraded from OpenStack Platform 9 to 10 and now aim to upgrade to OpenStack Platform 11.

BZ#1455793

> OpenStack Compute (nova) provides both versioned and unversioned notifications in RabbitMQ. However, due to the lack of consumers for versioned notifications, the versioned notifications queue grows quickly and causes RabbitMQ failures. This can hinder Compute operations such as instance creation and flavor creation. Red Hat is currently implementing fixes for RabbitMQ and director:
>
> https://bugzilla.redhat.com/show_bug.cgi?id=1478274
> https://bugzilla.redhat.com/show_bug.cgi?id=1488499
>
> The following article provides a workaround until Red Hat releases patches for this issue:
>
> https://access.redhat.com/solutions/3139721

BZ#1463058

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

BZ#1467849

> Previously, during the Red Hat OpenStack Platform 11 deployment, there is a race condition where, on occasion the ceilometer-upgrade runs at the same time as Apache is restarted due to other services being configured in the same step. This resulted in the ceilometer-upgrade to fail as it could not authenticate with the Identity service, as Apache was still not in the active state and aborting the deployment as a failure.
>
> With this update, as a workaround, when this happens you need to restart the overcloud deploy from where it fails and the deploy should get past this race condition and proceed with the deployment as normal. As a result, the deployment should be successful instead of failing with an error.

BZ#1488369

> RHEL overcloud images contain tuned version 2.8.
> In OVS-DPDK and SR-IOV deployments, tuned install and activation is done through the first-boot mechanism.
>
> This install and activation fails, as described in

> https://bugzilla.redhat.com/show_bug.cgi?id=1488369#c1
>
> You need to reboot the compute node to enforce the tuned profile.

3.2.4. Deprecated Functionality

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

BZ#1461990

> As of this release, Glance API V1 is no longer supported or available.

BZ#1488633

> The store_events option in undercloud.conf has been deprecated and is no longer supported. This option has been removed from the configuration.

CHAPTER 4. TECHNICAL NOTES

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

4.1. RHEA-2016:1245 — RED HAT OPENSTACK PLATFORM 11.0 BUG FIX AND ENHANCEMENT ADVISORY

The bugs contained in this section are addressed by advisory RHEA-2017:1245-02. Further information about this advisory is available at https://access.redhat.com/errata/RHEA-2017:1245.html.

instack-undercloud

BZ#1418828

> The undercloud stack rc file is a Keystone v2 rc. Previously, when switching from a v3 rc file (such as the v3 overcloudrc), some of the v3 environment variables would still be present. As a result, Keystone authentication may not work correctly.
>
> With this release, all OpenStack-related environment variables are cleared in stackrc before the undercloud values are set. As a result, variables from a previous rc file can no longer be present in the environment after sourcing stackrc, so Keystone authentication works correctly.

BZ#1256912

> The image_path parameter is no longer used. This update removes it from the undercloud configuration file.

BZ#1268451

> In certain situations, the undercloud virtual IPs were not correctly validated due to an error in the validation logic. Consequently, the undercloud could be deployed with incorrect virtual IPs. The error has been fixed. Now the virtual IPs are correctly validated. Any problem in virtual IP configuration is discovered before actual deployment.

openstack-cinder

BZ#1434494

> Previously, concurrent requests to create a volume from the same image could result in multiple entries in the Block Storage service's image cache. This resulted in duplicated image cache entries for the same image, which wasted space.
>
> This update adds a synchronization lock to prevent this. The first request to create a volume from an image will be cached, and all other requests will use the cached image.

openstack-glance

BZ#1396794

> With this enhancement, `glance-manage db purge` can now remove rows that are less than one day old. This was added because operators may need to run this operation on a regular basis.
> As a result, the value of the `age_in_days` option can be set to `0`.

openstack-gnocchi

BZ#1197163

> The Time Series Database as a Service (gnocchi) and Aodh API endpoints now expose a `/healthcheck` HTTP endpoint on the REST API. Requesting this endpoint allows you to check the status of the service, and does not require authentication.

openstack-heat

BZ#1414779

> Previously, when a pre-update hook was set on a resource that was in a FAILED state, the Orchestration service recorded an event indicating the hook was active. The service would then immediately create a replacement resource without waiting for the hook to be cleared by the user. As a result, the tripleoclient service believed the hook to be pending (based on the event), but fail upon trying to clear it as the replacement resource did not have a hook set. This, in turn, prevented the director from completing an overcloud update with the following message:
>
> ```
>     ERROR: The "pre-update" hook is not defined on SoftwareDeployment
>     "UpdateDeployment"
> ```
>
> This also affected other client-side applications that used hooks. In the director, this could have also resulted in UpdateDeployment executing on two Controller nodes simultaneously, instead of serialized so that only one Controller is updated at a time.
>
> With this release, the Orchestration service now pauses until the hook is cleared by the user, regardless of the state of the resource. This allows director overcloud updates to complete even when there is an UpdateDeployment resource in a FAILED state.

BZ#1320771

> Previously, while the Orchestration service could reset the status of resource when the state of the stack was incorrect, the service failed to do so when an update was retriggered. This resulted in resources being stuck in progress, which required database fixes to unblock the deployment.

> With this release, the Orchestration service now sets the status of all resources when it sets the status of the stack. This prevents the resources from getting stuck in progress, allowing operations to be retried successfully.

openstack-manila

BZ#1386249

> This update provides enhancements to the CephFS Native Driver in conjunction with the core OpenStack File Share Service (manila) infrastructure. The CephFS Native driver now supports read-only shares and improves recovery mode by deleting backend rules not in the 'access_list'.

openstack-manila-ui

BZ#1393893

> With this enhancement, you can now enable the creation of non-public shares in the Dashboard.
> You can configure Dashboard to hide the checkbox that enables users to mark shares as public during the creation process. The default option will be to create shares as private without the box checked.

openstack-neutron

BZ#1385338

> To implement the security groups trunk feature with neutron-openvswitch-agent, openvswitch firewall driver is required. This driver currently contains a bug 1444368 where ingress traffic is wrongly matched if there are two ports with same MAC address on different network segment on the same compute node.
>
> As a result, if a subport has the same MAC address as its parent port, ingress traffic won't be matched correctly for one of the ports.
>
> A workaround to achieve correctly handled traffic is to disable port-security on the parent port and subports.
>
> For example, to disable port security on port with UUID 12345, you need to remove security groups associated with the port:
> openstack port set --no-security-group --disable-port-security 12345
>
> Note that no security groups rules will be applied to that port and traffic will not be filtered or protected against ip/mac/arp spoofing.

BZ#1436576

> On DVR setups, the 'test_add_list_remove_router_on_l3_agent' from the

> 'test_l3_agent_scheduler.py' would not finish successfully. The testing procedure tried to bind a network interface to an L3 agent, although the interface had been bound to one previously, when a new router was created.
>
> The problem has been fixed. Now the interface will not be added to the router and assigned to the L3 agent until the test does so. As a result, the test finishes successfully.

openstack-neutron-lbaas

BZ#1325861

> This enhancement adds the ability to automatically reschedule load balancers from LBaaS agents that the server detects as dead. Previously, load balancers could be scheduled and realized across multiple LBaaS agents, however if a hypervisor died, the load balancers scheduled to that node would cease operation. With this update, these load balancers will be automatically rescheduled to a different agent. This feature is disabled by default and managed using `allow_automatic_lbaas_agent_failover`.

BZ#1326224

> This enahncement implements the 'ProcessMonitor' class in the 'HaproxyNSDriver' class (v2), This class utilizes the 'external_process' module to monitor and respawn HAProxy processes if and when needed. The LBaaS agent (v2) loads 'external_process' related options and take a configured action when HAProxy dies unexpectedly.

openstack-nova

BZ#1352922

> This release adds pagination support to avoid resource-consuming usage requests on systems with a large number of instances. The v2.40 microversion of the nova API simple-tenant-usage endpoints use new optional query parameters 'limit' and 'marker' for pagination. The 'marker' option sets the starting point and the 'limit' option sets the number of records to be displayed after the starting point. If 'limit' is not set, nova will use the configurable 'max_limit' (1000 by default). Although older microversions will not accept these new query parameters, they will start to enforce the max_limit and results may be truncated as a result. Consider using the new microversion to avoid DoS-like usage requests and potentially truncated responses.

openstack-sahara

BZ#1337664

> This update adds support for the version 5.1.0 MapR plugin.

openstack-selinux

BZ#1431556

> Because SELinux policies concerning launching instances with DPDK enabled are incomplete, launching instances using DPDK with SELinux in enforcing mode will cause the launch to fail and AVC denials will appear in /var/log/audit/audit.log* concerning openvswitch and svirt.
>
> As a workaround, set SELinux to permissive on each compute node where DPDK is utilized as documented in section 4.4.1.2 here:
>
> Permanent Changes in SELinux States and Modes
>
> This will allow DPDK-enabled virtual machines to launch. This is a workaround and is expected to be temporary while the issue is investigated further.

openstack-tripleo-common

BZ#1326549

> When deleting a node in heat, the deletion command finished and the prompt returned, although the process was still going on in the background. If another command followed immediately, a conflict would occur and the consequent command would fail. The behavior of the process has been changed. Now, the prompt will only return, when the process finishes completely.

BZ#1242422

> Automatic fencing setup can be used in director for easier High Availability deployments and upgrades. To benefit from the new feature, use the 'overcloud generate fencing' command.

openstack-tripleo-heat-templates

BZ#1425507

> If stopping the neutron-openvswitch-agent service, the stopping process sometimes took too long to exit gracefully and was killed by systemd. In this case, a running neutron-rootwrap-daemon remained in the system, which prevented the neutron-openvswitch-agent service to restart.
> The problem has been fixed. Now, an rpm scriplet detects the orphaned neutron-rootwrap-daemon and terminates it. As a result, the neutron-openvswitch-agent service starts and restarts successfully.

BZ#1435271

> With this release, 'clustercheck' will only run on nodes specified in the 'wsrep_cluster_address' option of Galera. This change was implemented to take into account use cases where Galera is run on a

> dedicated node (as is made possible with composable roles). Previously, during minor updates 'clustercheck' ran on all nodes running pacemaker, assuming Galera was also on the same node.

BZ#1312962

> The director set the 'tcp_list_options' stanza twice in '/etc/rabbitmq/rabbitmq.config'. This caused no adverse effects but could cause confusion in the future. This fix removes the redundant stanza. Only one 'tcp_list_options' stanza now appears in the configuration file.

BZ#1372589

> It is now possible to use puppet hieradata to set the max_files and max_processes for QEMU instances spawned by libvirtd. This can be done through an environment file containing the appropriate puppet classes. For example, to set the max_files and max_processes to 32768 and 131072 respectively, use:
>
> ```
> parameter_defaults:
>   ExtraConfig
>     nova::compute::libvirt::qemu::max_files: 32768
>     nova::compute::libvirt::qemu::max_processes: 131072
> ```
>
> This update also sets these values as the default, since QEMU instances launched by libvirtd might consume a large number of file descriptors or threads. This depends on Compute guest hosted on each compute node and of Ceph RBD images each instance attaches to. It is necessary to be able to configure these limits in large clusters.
>
> With these new default values, the Compute service should be able to use more than 700 OSDs. This was previously identified as the limit imposed by the low number of max_files (originally 1024).

BZ#1438890

> OpenStack Platform 10 included a broken Big Switch agent configuration. Deploying Big Switch agents with the provided heat templates resulted in deployment failures. This fix updates the heat templates to properly deploy Big Switch agents. Now the director correctly deploys the Big Switch agent service in composable roles.

BZ#1400262

> The default memory configuration for Memcached was 95 per cent of total available RAM, which could lead to resource contention. This fix lowers the default value to 50 per cent of total available RAM. You also can now configure this value using the 'MemcachedMaxMemory' setting. This helps reduce possible resource conflicts.

BZ#1440213

```
A bug in the overcloud package update script caused cluster services to
always restart even if no packages were available for update. This fix
corrects the check that determines if there are pending package updates.
If no packages updates are available, the yum update script exits and
does not restart cluster services.
```

BZ#1225069

```
For security reasons, the Overcloud only allows SSH key-based access by
default. You can set a root password on the disk image for the overcloud
using the virt-customize tool, which is found in the Red Hat Enterprise
Linux Extras channel. After installing the tool and downloading the
Overcloud images, use the following command to change the root password:

$ virt-customize -a overcloud-full.qcow2 --root-password
password:my_root_password

Perform this operation prior to uploading the images into glance with
the "openstack overcloud image upload" command.
```

openstack-tripleo-puppet-elements

BZ#1441923

```
The 'tuned-profiles-cpu-partitioning' package is now pre-installed in
the 'overcloud-full.qcow2' image. For DPDK deployments, this package is
necessary to help tune hosts and isolate CPU usage. The director
contains appropriate firstboot scripts to enable the 'tuned' service
with the necessary arguments.
```

os-net-config

BZ#1409097

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

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

puppet-ceph

BZ#1388515

```
When upgrading or deploying a Red Hat OpenStack Platform environment
integrated with an external Ceph Storage Cluster from an earlier version
(that is, Red Hat Ceph Storage 1.3), it is necessary to enable backwards
compatibility. To do so uncomment the following line in
```

```
environments/puppet-ceph-external.yaml during upgrade or deployment:

parameter_defaults:
  # Uncomment if connecting to a pre-Jewel or RHCS1.3 Ceph Cluster
  RbdDefaultFeatures: 1
```

BZ#1413980

> This release features the necessary puppet modules for deploying CephFS. This allows you to deploy the OpenStack Shared File System service (openstack-manila) with a CephFS back-end through the director.

puppet-pacemaker

BZ#1437417

> Previously, sometimes a deployment failed with the following error:
> Error: /Stage[main]/Pacemaker::Corosync/Exec[Start Cluster tripleo_cluster]/returns: change from notrun to 0 failed: /sbin/pcs cluster start --all returned 1 instead of one of 0
>
> With this update, a small race condition where puppet pacemaker could fail during cluster setup was closed. As a result, the deployment works correctly without errors.

BZ#1379741

> Previously, all pacemaker services had to be part of the same role.
>
> With this update, a new feature allows you to use composable roles with pacemaker managed services. This feature is needed in order to scale out pacemaker managed services on more and different nodes.

puppet-tripleo

BZ#1438602

> Previously, the OpenStack Dashboard service was configured in the wrong step of the deployment, resulting in horizon being temporarily unavailable during deployments and leading to additional 'httpd' service restarts.
>
> With this update, the OpenStack Dashboard configuration is fixed to occur at the same time as the rest of the 'httpd' configuration. As a result, horizon does not become temporarily unavailable when running the overcloud.

python-django-horizon

BZ#1271019

> The administrator needs to record the user credentials during the volume transfer operation and doing so by hand is inconvenient.
>
> With this update, a new button to download the credentials has been added to the volume transfer screen for saving the information easily. This allows the administrators to download and save a CSV file locally on their computer with the click of a button.

BZ#1388171

> To avoid memory bloat issues in the nova-api workers, pagination logic has been added to the simple-tenant-usage API extension.

BZ#1434704

> Previously, improper handling of the user IDs containing underscore in the code made it impossible to update project/domain members when the user IDs contained underscores.
>
> With this update, the code that handles the user IDs has been corrected to properly handle underscores. As a result, the project/domain members can now be updated even if they contain underscores.

python-heatclient

BZ#1437334

> After the optimization of event retrieval process, the 'openstack stack hook poll' command stopped returning pending hooks, even if they existed and should be returned. The problem was fixed. Now pending hooks are returned correctly.

python-openstackclient

BZ#1402772

> The '--os-interface' switch was ignored by 'openstack network' commands. Consequently, all such commands used the 'public' endpoint, although other interfaces were specified. The support for the switch has been added. Now the 'openstack network' commands correctly use the endpoint specified in '--os-interface' switch.

python-oslo-messaging

BZ#1427792

> The Remote Procedure Call (RPC) message acknowledgement in Oslo Messaging was not thread-safe. Consequently, a race condition caused an RPC timeout in Ceilometer. The message acknowledgement in Oslo Messaging

> has been fixed. Now, Ceilometer responds correctly.

BZ#1414497

> The Oslo Messaging did not initialize its configuration properly. As a result, the 'nova-manage' client failed during startup. The error has been fixed. Now 'nova-manage' starts correctly.

python-tripleoclient

BZ#1353049

> Previously, a failed update or upgrade would return an exit value of 0, so it was not possible to test for success based upon this value. With this update, a failed update or upgrade will throw an exception to signify to OpenStackClient that there is an error condition. As a result, OpenStackClient will only return an exit value of 0 on success, and a non-zero value after an error.

BZ#1400386

> The 'openstack overcloud image upload' ignored the '--image-path' argument when uploading or updating overcloud images. Consequently, only images in the working directory could be used. The support for the '--image-path' argument has been added. Now images from different directories, specified by the argument, can be uploaded flawlessly.

rhosp-director

BZ#1247019

> pacemaker continuously crashes when the fencing device name and the host name are the same. To avoid this problem, add the "fence-" prefix or the "-fence" suffix to the name of the fencing device. With the names configured like this, the cluster works without errors.