



Red Hat OpenStack Platform 12

Release Notes

Release details for Red Hat OpenStack Platform 12

Red Hat OpenStack Platform 12 Release Notes

Release details for Red Hat OpenStack Platform 12

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](https://creativecommons.org/licenses/by-sa/3.0/). 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, the Red Hat 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 is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

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

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

CHAPTER 1. INTRODUCTION	3
1.1. ABOUT THIS RELEASE	3
1.2. REQUIREMENTS	3
1.3. DEPLOYMENT LIMITS	3
1.4. DATABASE SIZE MANAGEMENT	4
1.5. CERTIFIED DRIVERS AND PLUG-INS	4
1.6. CERTIFIED GUEST OPERATING SYSTEMS	4
1.7. BARE METAL PROVISIONING SUPPORTED OPERATING SYSTEMS	4
1.8. HYPERVISOR SUPPORT	4
1.9. CONTENT DELIVERY NETWORK (CDN) CHANNELS	4
1.10. PRODUCT SUPPORT	6
1.11. KEY CHANGES TO THE DOCUMENTATION SET	7
CHAPTER 2. TOP NEW FEATURES	9
2.1. RED HAT OPENSTACK PLATFORM DIRECTOR	9
2.2. CONTAINERS	9
2.3. BARE METAL SERVICE	10
2.4. BLOCK STORAGE	11
2.5. CEPH STORAGE	11
2.6. COMPUTE	11
2.7. HIGH AVAILABILITY	12
2.8. IDENTITY	12
2.9. NETWORK FUNCTIONS VIRTUALIZATION	13
2.10. OBJECT STORAGE	13
2.11. OPENDAYLIGHT (TECHNOLOGY PREVIEW)	13
2.12. OPENSTACK DATA PROCESSING SERVICE	15
2.13. OPENSTACK NETWORKING	15
2.14. OPERATIONS TOOLING	16
2.15. SHARED FILE SYSTEM	16
2.16. TELEMETRY	16
2.17. TECHNOLOGY PREVIEWS	18
CHAPTER 3. RELEASE INFORMATION	22
3.1. RED HAT OPENSTACK PLATFORM 12 GA	22
3.2. RED HAT OPENSTACK PLATFORM 12 MAINTENANCE RELEASE JANUARY 2018	28
3.3. RED HAT OPENSTACK PLATFORM 12 MAINTENANCE RELEASE 28 MARCH 2018	29
3.4. RED HAT OPENSTACK PLATFORM 12 MAINTENANCE RELEASE 20 AUGUST 2018	30
3.5. RED HAT OPENSTACK PLATFORM 12 MAINTENANCE RELEASE 5 DECEMBER 2018	32
CHAPTER 4. TECHNICAL NOTES	34
4.1. RHEA-2017:3462 – RED HAT OPENSTACK PLATFORM 12.0 ENHANCEMENT ADVISORY	34
4.2. RHEA-2018:2331 – RED HAT OPENSTACK PLATFORM 12.0 ENHANCEMENT ADVISORY AUGUST 2018	40
4.3. RHEA-2018:2332 – RED HAT OPENSTACK PLATFORM 12.0 SECURITY ADVISORY AUGUST 2018	46

CHAPTER 1. INTRODUCTION

1.1. ABOUT THIS RELEASE

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

<https://releases.openstack.org/pike/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 7.4.

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 12.

You can install Red Hat OpenStack Platform 12 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 (x86_64)

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 12 for RHEL 7 (RPMs)	rhel-7-server-openstack-12-rpms
Red Hat Enterprise Linux 7 Server - Extras (RPMs)	rhel-7-server-extras-rpms

Table 1.2. Optional Channels (x86_64)

Channel	Repository Name
Red Hat Enterprise Linux 7 Server - Optional	rhel-7-server-optional-rpms
Red Hat OpenStack Platform 12 Operational Tools for RHEL 7 (RPMs)	rhel-7-server-openstack-12-optools-rpms

Table 1.3. Required Channels (ppc64le)

Channel	Repository Name
Red Hat Enterprise Linux for IBM Power, little endian	rhel-7-for-power-le-rpms

Channel	Repository Name
Red Hat OpenStack Platform 12 for RHEL 7 (RPMs)	rhel-7-server-openstack-12-for-power-le-rpms

Channels to Disable

The following table outlines the channels you must disable to ensure Red Hat OpenStack Platform 12 functions correctly.

Table 1.4. 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>.

1.11. KEY CHANGES TO THE DOCUMENTATION SET

Several notable changes were made to the Red Hat OpenStack Platform documentation set for this release. The following list outlines and explains these changes.

Configuration Reference and Command-Line Interface Reference

The *Configuration Reference* and *Command-Line Interface Reference* documents are not available with the general availability of Red Hat OpenStack Platform 12.

Both documents have a dependency on source content generated by openstack.org. In the Pike release, the format and location of this content changed. As a result, the scope of the work required extended beyond the Red Hat OpenStack Platform 12 GA schedule. The documents will be compiled and published as an asynchronous release.

Manual Installation Procedures

The *Manual Installation Procedures* document has been removed from the documentation set starting from Red Hat OpenStack Platform 12.

Red Hat supports installation processes performed only using the Red Hat OpenStack Platform director and the official documentation steps.

Users can find manual installation information, for their reference, on the OpenStack website: <https://docs.openstack.org/pike/install>. These procedures will not be supported by Red Hat.

Manual Upgrades

Red Hat OpenStack Platform 12 does not support manual upgrade steps performed without director, and, as a result, no documentation is available for this scenario. For supported upgrade scenarios, using Red Hat OpenStack Platform director, see [Upgrading Red Hat OpenStack Platform](#).

OpenStack Benchmarking Service Guide

The *OpenStack Benchmarking Service* document was outdated and contained incorrect information. It has been removed from the documentation set across all versions. The following Bugzilla ticket requests a full review of the document: https://bugzilla.redhat.com/show_bug.cgi?id=1459469.

Red Hat Ceph Storage for the Overcloud

The *Red Hat Ceph Storage for the Overcloud* document has been replaced by two new guides, which describe two available options for working with Red Hat Ceph Storage in the overcloud: [Integrating an Overcloud with an Existing Red Hat Ceph Cluster](#) and [Deploying an Overcloud with Containerized Red Hat Ceph](#).

RPM-Based Overcloud Installation

RPM packages for Red Hat OpenStack Platform 12 are still shipped alongside container images; however, the official installer provided by Red Hat, Red Hat OpenStack Platform director, does not

support deployments of non-containerized (RPM-based) Red Hat OpenStack Platform 12. As a result, instructions for deploying an RPM-based overcloud are not provided in the documentation. For supported deployment scenarios, see [Director Installation and Usage](#).

Although installation procedures are not provided or supported for RPM-based deployments, environments resulting from manual deployment are still supported if they comply with Red Hat support policy: <https://access.redhat.com/articles/2477851>.

VMware Integration Guide

The *VMware Integration Guide* has been removed from the documentation set across all versions. The integration described in the document is no longer supported.

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.

Prompt Changes

Sourcing a settings file on the undercloud, such as **stackrc** or **overcloudrc**, changes Prompt String 1 (PS1) to include the cloud name. This helps identify the cloud currently being accessed. For example, if you source the **stackrc** file, the prompt appears with an **(undercloud)** prefix:

```
[stack@director-12 ~]$ source ~/stackrc  
(undercloud) [stack@director-12 ~]$
```

Registration through a HTTP Proxy

The director provides updated templates to register your overcloud through a HTTP proxy.

New Custom Roles Generation

The director provides the ability to create a **roles_data** file from individual custom role files. This simplifies the management of individual custom roles. The director also includes a default set of role files to help you get started.

Node Blacklist

The director now accepts a node blacklist using the **DeploymentServerBlacklist** parameter. This parameter isolates a list of nodes from receiving updated parameters and resources during the execution of **openstack overcloud deploy**. This parameter is useful to scale additional nodes while the existing nodes remain untouched during the deployment process.

Composable Networks

Previously, the Bare Metal service could only use networks defined in the director templates. It is now possible to compose custom networks for the director to create during an overcloud deployment or update. You can also now assign custom labels to the director's template-defined networks.

UI: Improved Node Management

The director's web UI now provides more detail for each node and additional functions to manage nodes. You can view this additional information on the **Nodes** screen of the director's web UI.

UI: Improved Role Assignment

The director's web UI includes a simplified role assignment for nodes. The UI uses a spinner to automatically assign a selected number of nodes per role. You can also manually assign specific nodes to roles in the **Nodes** screen of the director's web UI.

2.2. CONTAINERS

This section outlines the introduction of containerization in Red Hat OpenStack Platform.

Containerized Overcloud

The Red Hat OpenStack Platform director now creates an overcloud consisting of containerized services. Users can implement the following sources for their container images:

- Remote from **registry.access.redhat.com**
- Locally from the undercloud (images initially pulled from **registry.access.redhat.com**)
- Red Hat Satellite 6 (images synchronized from **registry.access.redhat.com**)

The overcloud continues to support composable service infrastructure using containers to augment existing services. Note that the only services not containerized by default are:

- OpenStack Networking (neutron)
- OpenStack Block Storage (cinder)
- OpenStack Shared File Systems (Manila)

Red Hat provides the containers for these services as a **technical preview only**.

Containerized Upgrades

The Red Hat OpenStack Platform director provides an upgrade path from a non-containerized Red Hat OpenStack Platform 11 overcloud to containerized Red Hat OpenStack Platform 12 overcloud.

2.3. BARE METAL SERVICE

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

L3 Routed Spine/Leaf Network Topology

With spine/leaf, the bare metal network now uses layer 3 routing. This new topology makes full use of connections through equal-cost multipathing (ECMP). You can use the new topology only with the Compute and Ceph storage roles. It is not yet possible to use this routing for the provisioning network.

Node Auto-Discovery

Previously, writing an instack.json file was the only way to add overcloud nodes in bulk. The Bare Metal Service can now discover unidentified nodes automatically, without an instack.json file.

Redfish Support

The Redfish API is an open standard for the management of hardware. The Bare Metal service now includes the Redfish API driver. To manage servers compliant with the Redfish protocol, set the driver property to **redfish**.

Whole-disk Overcloud image support

The Bare Metal Service now supports whole-disk images for the overcloud. Previously, initrd and vmlinuz images were required in addition to the qcow2 image. Now, the Bare Metal Service can accept a single qcow2 image upload as a full disk image. You must build the whole-disk image before you deploy it.

2.4. BLOCK STORAGE

This section outlines the top new features for Block Storage.

Capacity-derived QoS Limits

Users can now use volume types to set deterministic IOPS throughput based on the size of provisioned volumes. This simplifies how storage resources are allocated to users -- namely, through pre-determined (and, ultimately, highly predictable) throughput rates based on requested volume size.

Veritas HyperScale Support

The Block Storage service now supports the HyperScale driver. HyperScale is a software-defined storage solution that uses a dual-plane architecture to decouple storage management tasks from workload processing at the Compute plane. This technique helps make efficient use of storage directly attached to Compute nodes, thereby minimizing total cost of ownership without compromising performance.

Veritas HyperScale requires binaries, puppet modules, and Heat templates provided directly by Veritas. For an overview, see [HyperScale for OpenStack](#); for deployment and usage documentation, see [HyperScale for OpenStack guides for Linux](#).

2.5. CEPH STORAGE

This section outlines the top new features for Ceph Storage.

Containerized Ceph Deployment

The director can now deploy a containerized Red Hat Ceph cluster. To do this, the director uses built-in heat templates and environment files that work with Ansible playbooks available through the **ceph-ansible** project.

Improved Resource Management for HCI

Previously, when deploying Hyper-Converged Infrastructure (HCI) users had to manually configure resource isolation on hyper-converged Compute nodes. The director can now use OpenStack Workflow to derive HCI-suitable CPU and RAM allocation settings and apply them.

2.6. COMPUTE

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

Emulator Thread Policies

The Compute scheduler determines the CPU resource utilization and places instances based on the number of virtual CPUs (vCPUs) in the flavor. There are a number of hypervisor operations that are performed on the host on behalf of the guest instance. The **libvirt** driver implements a generic placement policy for KVM, which allows QEMU emulator threads to float across the same physical CPUs (pCPUs) that the vCPUs are running on. This leads to the emulator threads using time borrowed from the vCPU operations.

With this release, Compute reserves one vCPU for running non-realtime workloads using the **hw:emulator_threads_policy=isolate** option. Before you enable the emulator threads placement policy on an instance flavor, you must set the **hw:cpu_policy** option to *dedicated*.

Reserve (weight) SR-IOV Capable NUMA Nodes

This release includes updates to the filter scheduler and resource tracker to place non-PCI instances on non-PCI NUMA nodes. Instances not bound to PCI devices will be preferably placed on hosts without PCI devices. If there is no host without PCI devices, then hosts with PCI devices will be used. To enable this option, use the new **nova.scheduler.weights.all_weighers** PCI weigher option. You can also enable this option manually using the **filter_scheduler.weight_classes** configuration option.

2.7. HIGH AVAILABILITY

This section outlines the top new features for high availability.

Containerized High Availability Reference Architecture

The Instance High Availability (Instance HA) reference architecture is now provided in containers that you can deploy on a Red Hat Enterprise Linux Atomic Host with the Red Hat OpenStack director. The Instance HA configuration is provided in an agent container, which then deploys the application containers and shared services across the cluster.

The following Instance HA components and managed services are now delivered as containers:

- Pacemaker
- Pacemaker_remote
- Corosync
- Ancillary supporting components
- Galera (MariaDB)
- RabbitMQ
- HAProxy
- Cinder-backup
- Cinder-volume
- Manila-share
- Redis
- Virtual-ips
- memcached

Health Checks with httpchk

Instance HA now uses the httpchk to check the health of compatible service nodes in the cluster.

2.8. IDENTITY

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

Novajoin Availability for Infrastructure

Novajoin allows you to enroll undercloud and overcloud nodes with Red Hat Identity Management (IdM). As a result, you can now use IdM features with your OpenStack deployment, including identities, kerberos credentials, and access controls.

TLS Coverage

Red Hat OpenStack Platform 12 includes TLS support for MariaDB, RabbitMQ and internal services endpoints.

2.9. NETWORK FUNCTIONS VIRTUALIZATION

This section outlines the top new features for Network Functions Virtualization (NFV).

Easy Heterogeneous Cluster Management

Easy heterogeneous cluster management allows you to set different service parameter values to match varying node capabilities or tuning needs. For example, if node 1 had more RAM than node 2, you could not previously take advantage of the added RAM with two different Compute roles, since the service parameters were defined globally. You can now combine composable roles with role-specific parameters to define unique parameters that match the capabilities of different nodes or different tuning needs.

OpenDaylight (Technology Preview)

The OpenDaylight software-defined networking (SDN) controller is now integrated into Red Hat OpenStack Platform.

OVS-DPDK Ease of Deployment

The Red Hat OpenStack Platform simplifies OVS-DPDK deployments through a predefined Mistral workflow to auto generate OVS-DPDK parameters. You now need to decide on two simple parameters (the minimum number of CPU threads used for DPDK PMD, and the percent of available memory to reserve for Hugepages). Based on this information and the hardware introspection results of your bare-metal nodes, the workflow calculates the remaining eight OVS-DPDK parameters needed for your deployment.

NUMA Topology through Bare Metal Introspection

To ease deployment, you can now retrieve the NUMA topology details from Compute nodes with the Bare Metal hardware inspection service. The retrieved NUMA topology details include NUMA nodes, associated RAM, NICs, and physical CPU cores with sibling pairs.

2.10. OBJECT STORAGE

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

Stand-Alone Object Storage Deployments

With this release, users can now configure a new overcloud deployment to use an existing Object Storage cluster.

2.11. OPENDAYLIGHT (TECHNOLOGY PREVIEW)

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

Improved Red Hat OpenStack Platform director integration

The Red Hat OpenStack Platform director installs and manages a complete OpenStack environment. With Red Hat OpenStack Platform 12, the director can deploy and configure OpenStack to work with OpenDaylight. OpenDaylight can run together with the OpenStack overcloud controller role, or in a separate custom role on a different node.

In Red Hat OpenStack Platform 12, OpenDaylight is installed and run in containers. This provides more flexibility to its maintenance and use.

IPv6

OpenDaylight in Red Hat OpenStack Platform 12 brings some feature parity in IPv6 use-cases with OpenStack neutron ML2/OVS implementation. These use-cases include:

- IPv6 addressing support including SLAAC
- Stateless and Stateful DHCPv6
- IPv6 Security Groups with allowed address pairs
- IPv6 communication among virtual machines in the same network
- IPv6 East-West routing support

VLAN-aware virtual machines

VLAN-aware virtual machines (or virtual machines with trunking support) allow an instance to be connected to one or more networks over one virtual NIC (vNIC). Multiple networks can be presented to an instance by connecting it to a single port. Network trunking lets users create a port, associate it with a trunk, and launch an instance on that port. Later, additional networks can be attached to or detached from the instance dynamically without interrupting the operations of that instance.

SNAT

Red Hat OpenStack Platform 12 introduces the conntrack-based SNAT, where it uses OVS *netfilter* to maintain translations. A switch is selected as the NAPT switch per router, and does the centralized translation. All the other switches send the packet to centralized switch for SNAT. If a NAPT switch goes down, an alternate switch is selected for the translations and the existing translations will be lost on a failover.

SR-IOV Integration

OpenDaylight in Red Hat OpenStack Platform 12 can be deployed with compute nodes that support SR-IOV. It is also possible to create mixed environments with both OVS and SR-IOV nodes in a single OpenDaylight installation. The SR-IOV deployment requires the neutron SR-IOV agent in order to configure the virtual functions (VFs), which are directly passed to the compute instance when it is deployed as a network port.

Controller Clustering

The OpenDaylight Controller in Red Hat OpenStack Platform 12 supports a cluster-based High Availability model. Several instances of the OpenDaylight Controller form a *Controller Cluster*. Together, they work as one logical controller. The service provided by the controller (viewed as a logical unit) will continue to operate as long as a majority of the controller instances are functional and able to communicate with each other.

The Red Hat OpenDaylight Clustering model provides both High Availability and horizontal scaling: more nodes can be added to absorb more load, if necessary.

OVS-DPDK

OpenDaylight in Red Hat OpenStack Platform 12 may be deployed with Open vSwitch Data Plane Development Kit (DPDK) acceleration with director. This deployment offers higher data plane performance as packets are processed in user space rather than in the kernel.

L2GW/HW-VTEP

Red Hat OpenStack Platform 12 supports L2GW to integrate traditional bare-metal services into a neutron overlay. This is especially useful for bridging external physical workloads into a neutron tenant network, for bringing a bare metal server (managed by OpenStack) into a tenant network, and bridging SR-IOV traffic into a VXLAN overlay. This fact takes advantage of the line-rate speed of SR-IOV and the benefits of an overlay network to interconnect SR-IOV virtual machines.

The networking-odl Package

Red Hat OpenStack Platform 12 offers a new version of the networking-odl package that brings important changes. It introduces the **port status update support** command that provides accurate information on the port status and when the port is available for a virtual machine to use. The default port binding changes from network topology to *pseudo agent* based. The network topology binding support is not available in this release. Customers using their network topology based on port binding should migrate to pseudo agent based port binding (**pseudo-agentdb-binding**).

2.12. OPENSTACK DATA PROCESSING SERVICE

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

OpenStack Data Processing Service Integration With Baremetal to Tenant

The OpenStack Data Processing service can help improve performance by removing the hypervisor abstraction layer. The OpenStack Bare Metal Provisioning (ironic) service provides an API and a Compute driver to serve bare metal instances using the same Compute and OpenStack Networking APIs.

This release adds the support to install and configure the OpenStack Bare Metal Provisioning service and Compute to serve the bare metal instance to the tenant. For Red Hat OpenStack Platform deployments with both virtual and bare metal instances, you need to use host aggregates as follows:

- One for all the bare metal hosts
- One for all the virtual Compute nodes

Support for the Cloudera Distribution of Apache Hadoop (CDH) 5.11

You can deploy the Cloudera Distribution of Apache Hadoop (CDH) 5.11 on Red Hat OpenStack Platform. CDH can store, process, and analyze large and diverse data volumes with the latest big data processing techniques such as Spark and Impala.

2.13. OPENSTACK NETWORKING

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

Native Open vSwitch Firewall Driver

The OVS firewall driver has graduated from Technology Preview to full support. 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.

Layer-2 Gateway API

The Layer-2 Gateway is a service plugin which allows you to bridge networks together so they appear as a single L2 broadcast domain. This update introduces support for the Layer-2 Gateway API.

BGP/VPN API

OpenStack Networking now supports BGPVPN capabilities. BGPVPN allows your instances to connect to your existing layer 3 VPN services. Once a BGPVPN network is created, you can associate it with a project, allowing the project's users to connect to the BGPVPN network.

2.14. OPERATIONS TOOLING

This section outlines the top new features for operations tooling.

SSL Support in the Monitoring Agent

You can now configure the Monitoring Agent (Sensu client) to connect to the RabbitMQ instance with SSL. To do this, you define the SSL connection parameters and certificates in the monitoring environment YAML file.

Integration with Red Hat Enterprise Common Logging

You can now use the Red Hat Enterprise Common Logging solution to collect logs from Red Hat OpenStack Platform. To do this, you configure the Log Collection Agent (Fluentd) to send the log files to the central logging collector.

Containerized Monitoring and Logging Tools

Some monitoring and logging tools are now provided in containers that you can deploy on a Red Hat Enterprise Linux Atomic Host with the Red Hat OpenStack director.

The following operation tools are now delivered in containers:

- Availability monitoring (Sensu)
- Performance monitoring (Collectd)
- Log aggregation (Fluentd)

2.15. SHARED FILE SYSTEM

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

Telemetry Integration

The Shared File System service can now send usage statistics to the Telemetry service. This allows users to monitor events in the life-cycle of shares.

2.16. TELEMETRY

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

OpenStack Telemetry Metrics (gnocchi) at Scale

Telemetry used MongoDB and Telemetry API to store metrics, while the performance was acceptable when it came to storing the metrics, the usage is limited as you cannot retrieve and exploit the stored information.

The OpenStack Telemetry Metrics (gnocchi) service uses a new distributed selective acknowledgements (SACKs) mechanism and scheduling algorithm for the **gnocchi-metricd** daemon improving the performance at a larger scale. The default settings are enhanced to adapt to cloud deployment of larger sizes.

Intel Cache Monitoring Technology (CMT)

Cache Monitoring Technology (CMT) allows you to monitor cache-related statistics on an Intel platform. Telemetry now supports CMT reporting using the **collectd** daemon.

This release adds a new meter to collect the L3 cache usage statistics for each virtual machine. You can enable the **cmt** plugin with **LibvirtEnabledPerfEvents** parameter in **nova-libvirt.yaml** file.

Containerization of Telemetry Services

With this release, Red Hat OpenStack Platform can create a cloud that uses containers to host its services. Each service is isolated within its own container on the host node. Each container connects to and shares the host's own network. As a result, the host node exposes the API ports of each service on its own network. Telemetry service can now be hosted on containers. This makes upgrades easy.

OpenStack Telemetry Event Storage (panko) Deprecation

The OpenStack Telemetry Event Storage service is officially now deprecated. Support for panko will be limited to usage from Red Hat Cloudforms only. Red Hat does not recommend using panko outside of the Red Hat Cloudforms use-case. You can use the following options instead of using panko:

- Poll the OpenStack Telemetry Metrics (gnocchi) service instead of polling panko. This gives you access to the resource history.
- Use the OpenStack Telemetry Alarming (aodh) service to trigger alarms when an event arises. You can use OpenStack Messaging Service (zaqar) to store alarms in a queue if an application cannot be reached directly by the OpenStack Telemetry Alarming (aodh) service.

Telemetry API and ceilometer-collector Deprecation

The Telemetry API service is now deprecated. It is replaced by the OpenStack Telemetry Metrics (gnocchi) service, and the OpenStack Telemetry Alarming (aodh) service APIs. You should begin to switch to the Telemetry API service instead. In Red Hat OpenStack Platform 12, the Telemetry API is disabled by default, with the option to enable it only if required.

The **ceilometer-collector** service is deprecated. You can now use the **ceilometer-notification-agent** daemon because the Telemetry polling agent sends the messages from the sample file to the **ceilometer-notification-agent** daemon.

NOTE: Ceilometer as a whole is not deprecated, just the Telemetry API service and the **ceilometer-collector** service.

2.17. TECHNOLOGY PREVIEWS

This section outlines features that are in technology preview in Red Hat OpenStack Platform 12.



NOTE

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

2.17.1. New Technology Previews

The following new features are provided as technology previews:

Octavia LBaaS

Octavia is a new component that can be used as a back-end plug-in for the LBaaS v2 API and is intended to replace the current HAProxy-based implementation.

Open Virtual Network (OVN)

OVN is an Open vSwitch-based network virtualization solution for supplying network services to instances.

Red Hat OpenStack Platform for POWER

You can now deploy pre-provisioned overcloud Compute nodes on IBM POWER8 little endian hardware.

2.17.2. Previously Released Technology Previews

The following features remain 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.

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 backend 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.

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 as a technology preview as well.

DNS-as-a-Service (DNSaaS)

Red Hat OpenStack Platform 12 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 12 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 12 is limited to the modules required to support OpenStack deployments using NetVirt, and is based on the upstream Carbon version.

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

Real Time KVM Integration

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

Red Hat SSO

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

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 12 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#1117883

This update provides the Docker image for the Keystone service.

BZ#1276147

This update adds support to OpenStack Bare Metal (ironic) for the Emulex hardware iSCSI (be2iscsi) ramdisk.

BZ#1277652

This update adds new commands that allow you to determine host to IP mapping from the undercloud without needing to access the hosts directly.

You can show which IP addresses are assigned to which host and to which port with the following command: `openstack stack output show overcloud HostsEntry -c output_value -f value`

Use `grep` to filter the results for a specific host.

You can also map the hosts to bare metal nodes with the following command: `openstack baremetal node list --fields uuid name instance_info -f json`

BZ#1293435

Uploading to and downloading from Cinder volumes with Glance is now supported with the Cinder backend driver.

Note: This update does not include support for Ceph RBD. Use the Ceph backend driver to perform RBD operations on Ceph volumes.

BZ#1301549

The update adds a new validation to check the overcloud's network environment. This helps avoid any conflicts with IP addresses, VLANs, and allocation pool when deploying your overcloud.

BZ#1334545

You can now set QoS IOPS limits that scale per GB size of the volume with the options "total_iops_sec_per_gb", "read_iops_sec_per_gb", and "write_iops_sec_per_gb".

For example, if you set the total_iops_sec_per_gb=1000 option, you will get 1000 IOPS for a 1GB volume, 2000 IOPS for a 2GB volume, and so on.

BZ#1368512

The update adds a new validation to check the hardware resource on the undercloud before an deployment or upgrade. The validation ensures the undercloud meets the necessary disk space and memory requirements prior to a deployment or upgrade.

BZ#1383576

This update adds an action to "Manage Nodes" through the director UI. This action switches nodes to a "manageable" state so the director can perform introspection through the UI.

BZ#1406102

Director now supports the creation of custom networks during the deployment and update phases. These additional networks can be used for dedicated network controllers, Ironic baremetal nodes, system management, or to create separate networks for different roles.

A single data file ('network_data.yaml') manages the list of networks that will be deployed. The role definition process then assigns the networks to the required roles.

BZ#1430885

This update increases the granularity of the deployment progress bar. This is achieved with an increase in the nesting level that retrieves the stack resources. This provides more accurate progress of a deployment.

BZ#1434929

Previously, the OS_IMAGE_API_VERSION and the OS_VOLUME_API_VERSION environment variables were not set, which forced Glance and Cinder to fall back to the default API versions. For Cinder, this was the older v2 API.

With this update, the overcloudrc file now sets the environment variables to specify the API versions for Glance and Cinder.

3.1.2. Technology Preview

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

BZ#1300425

With the Manila service, you can now create shares within Consistency Groups to guarantee snapshot consistency across multiple shares. Driver vendors must report this capability and implement its functions to work according to the back end.

This feature is not recommended for production cloud environments, as it is still in its experimental stage.

BZ#1418433

Containerized deployment of the OpenStack File Share Service (manila) is available as a technology preview in this release. By default, Manila, Cinder, and Neutron will still be deployed on bare metal machines.

BZ#1513109

POWER-8 (ppc64le) Compute support is now available as a technology preview.

3.1.3. Release Notes

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

BZ#1463355

When TLS everywhere is enabled, the HAProxy stats interface will also use TLS. As a result, you will need to access the interface through the individual node's ctlplane address, which is either the actual IP address or the FQDN (using the convention <node name>.ctlplane.<domain>, for example, overcloud-controller-0.ctlplane.example.com). This setting can be configured by the ``CloudNameCtlplane`` parameter in ``tripleo-heat-templates``. Note that you can still use the ``haproxy_stats_certificate`` parameter from the HAProxy class, and it will take precedence if set.

3.1.4. Known Issues

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

BZ#1552234

There is currently a known issue where you cannot use ACLs to make a container public for anonymous access. This issue arises when sending ``POST`` operations to Swift that specify a ``**`` value in the ``X-Container-Read`` or ``X-Container-Write`` settings.

BZ#1321179

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

BZ#1384845

When an overcloud image is shipped with 'tuned' version lower than 2.7.1-4, you should apply a manual update of the 'tuned' package to the overcloud image. If the 'tuned' version is equal to 2.7.1-4 or higher, you should provide the list of the core to 'tuned' and activate the profile, for example:

```
# echo "isolated_cores=2,4,6,8,10,12,14,18,20,22,24,26,28,30" >> /etc/tuned/cpu-partitioning-variables.conf
# tuned-adm profile cpu-partitioning
```

This is a known issue until the 'tuned' packages are available in the Centos repositories.

BZ#1385347

The '--controller-count' option for the 'openstack overcloud deploy' command sets the 'NeutronDhcpAgentsPerNetwork' parameter. When deploying a custom Networker role that hosts the OpenStack Networking (neutron) DHCP Agent, the 'NeutronDhcpAgentsPerNetwork' parameter might not set to the correct value. As a workaround, set the 'NeutronDhcpAgentsPerNetwork' parameter manually using an environment file. For example:

```
----
parameter_defaults:
  NeutronDhcpAgentsPerNetwork: 3
----
```

This sets 'NeutronDhcpAgentsPerNetwork' to the correct value.

BZ#1486995

When using an NFS back end for the Image service (glance), attempting to create an image will fail with a permission error. This is because the user ID on the host and container differ, and also because puppet cannot mount the NFS endpoint successfully on the container.

BZ#1487920

Encrypted volumes cannot attach correctly to instances in containerized environments. The Compute service runs "cryptsetup luksOpen", which waits for the udev device creation process to finish. This process does not actually finish, which causes the command to hang.

Workaround: Restart the containerized Compute service with the docker option "--ipc=host".

BZ#1508438

For containerized OpenStack services, configuration files are now installed in each container. However, some OpenStack services are not containerized yet, and configuration files for those services are still installed on the bare metal nodes.

If you need to access or modify configuration files for containerized services, use `/var/log/config-data/<container name>/<config path>`. For services that are not containerized yet, use `/etc/<service>`.

BZ#1516911

In HP DL 360/380 Gen9, the DIMM format does not match the regex query.

In order to PASS on this, you must cherry-pick the HW patches in comment #2.

BZ#1519057

There is currently a known issue with LDAP integration for Red Hat OpenStack Platform. At present, the ``keystone_domain_config`` tag is missing from ``keystone.yaml``, preventing Puppet from properly applying the required configuration files. Consequently, LDAP integration with Red Hat OpenStack Platform will not be properly configured. As a workaround, you will need to manually edit ``keystone.yaml`` and add the missing tag. There are two ways to do this:

1. Edit the the file directly:

- a. Log into the undercloud as the stack user.
- b. Open the `keystone.yaml` in the editor of your choice. For example:


```
`sudo vi /usr/share/openstack-tripleo-heat-templates/docker/services/keystone.yaml`
```
- c. Append the missing puppet tag, ``keystone_domain_config``, to line 94. For example:


```
`puppet_tags: keystone_config`
```

Changes to:

```
`puppet_tags: keystone_config,keystone_domain_config`
```

- d. Save and close ``keystone.yaml``.
- e. Verify you see the missing tag in the ``keystone.yaml`` file. The following command should return `'1'`:

```
`cat /usr/share/openstack-tripleo-heat-templates/docker/sercies/keystone.yaml | grep 'puppet_tags: keystone_config,keystone_domain_config' | wc -l`
```

2. Or, use `sed` to edit the file inline:

- a. Login to the undercloud as the stack user.
- b. Run the following command to add the missing puppet tag:

```
`sed -i 's/puppet_tags\[: keystone_config/puppet_tags\[: keystone_config,keystone_domain_config/' /usr/share/openstack-tripleo-heat-templates/docker/services/keystone.yaml`
```

- c. Verify you see the missing tag in the `keystone.yaml` file The following command should return `'1'`:

```
`cat /usr/share/openstack-tripleo-heat-templates/docker/sercies/keystone.yaml | grep 'puppet_tags: keystone_config,keystone_domain_config' | wc -l`
```

BZ#1519536

You must manually discover the latest Docker image tags for current container images that are stored in Red Hat Satellite. For more information, see the Red Hat Satellite documentation: https://access.redhat.com/documentation/en-us/red_hat_satellite/6.2/html/content_management_guide/managing_container_images#managing_container_images_with_docker_tags

BZ#1520004

It is only possible to deploy Ceph storage servers if their disk devices are homogeneous.

BZ#1522872

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#1525520

For deployments using OVN as the ML2 mechanism driver, only nodes with connectivity to the external networks are eligible to schedule the router gateway ports on them. However, there is currently a known issue that will make all nodes as eligible, which becomes a problem when the Compute nodes do not have external connectivity. As a result, if a router gateway port is scheduled on Compute nodes without external connectivity, ingress and egress connections for the external networks will not work; in which case the router gateway port has to be rescheduled to a controller node. As a workaround, you can provide connectivity on all your compute nodes, or you can consider deleting NeutronBridgeMappings, or set it to datacentre:br-ex. For more information, see https://bugzilla.redhat.com/show_bug.cgi?id=1525520 and https://bugzilla.redhat.com/show_bug.cgi?id=1510879.

3.1.5. Deprecated Functionality

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

BZ#1417221

The Panko service is officially deprecated in OpenStack version 12. Support for panko will be limited to usage from cloudforms only. We do not recommend using panko outside of the cloudforms use case.

BZ#1427719

VPN-as-a-Service (VPNaaS) VPNaaS was deprecated in Red Hat OpenStack Platform 11, and has now been removed in Red Hat OpenStack Platform 12.

BZ#1489801

MongoDB is no longer used by Red Hat OpenStack Platform. Previously, it was used for Telemetry (which now uses Gnocchi) and Zaqr on the undercloud (which is moving to Redis). As a result, 'mongodb', 'puppet-mongodb', and 'v8' are no longer included.

BZ#1510716

File injection from the Compute REST API. This will continue to be supported for now if using API microversion < 2.56. However, Compute will eventually remove this functionality. The changes are as follows:

- * Deprecate the 'personality' parameter from the 'POST /servers' create server API and the 'POST /servers/{server_id}/action' rebuild server API. Specifying the 'personality' parameter in the request body to either of these APIs will result in a '400 Bad Request' error response.
- * Add support to pass 'user_data' to the rebuild server API as a result of this change.
- * Stop returning 'maxPersonality' and 'maxPersonalitySize' response values from the 'GET /limits' API.
- * Stop accepting and returning 'injected_files', 'injected_file_path_bytes', 'injected_file_content_bytes' from the 'os-quota-sets' and 'os-quota-class-sets' APIs.
- * Removes Compute API extensions including server extensions, flavor extensions and image extensions.. The extensions code have their own policy and there is no option to enable or disable these extensions in the API, leading to interoperability issues.
- * Removes the 'hide_server_address_states' configuration option which allows you to configure the server states to hide the address and the hide server address policy. Also, removes the 'os_compute_api:os-hide-server-addresses' policy as it is no longer necessary.

3.2. RED HAT OPENSTACK PLATFORM 12 MAINTENANCE RELEASE JANUARY 2018

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

3.2.1. Known Issues

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

BZ#1552234

There is currently a known issue where you cannot use ACLs to make a container public for anonymous access. This issue arises when sending `POST` operations to Swift that specify a `**` value in the `X-Container-Read` or `X-Container-Write` settings.

BZ#1469434

When using the Docker CLI to report the state of running containers, the nova_migration_target container might be incorrectly reported as "unhealthy". This is due to an issue with the health check itself, and not with an accurate reflection of the state of the running container.

BZ#1519057

There is currently a known issue with LDAP integration for Red Hat OpenStack Platform. At present, the `keystone_domain_conf` tag is missing from `keystone.yaml`, preventing Puppet from properly applying the required configuration files. Consequently, LDAP integration with Red

Hat OpenStack Platform will not be properly configured. As a workaround, you will need to manually edit `keystone.yaml` and add the missing tag. There are two ways to do this:

1. Edit the the file directly:

- a. Log into the undercloud as the stack user.
- b. Open the `keystone.yaml` in the editor of your choice. For example:


```
sudo vi /usr/share/openstack-tripleo-heat-templates/docker/services/keystone.yaml
```
- c. Append the missing puppet tag, `keystone_domain_config`, to line 94. For example:


```
puppet_tags: keystone_config
```

 Changes to:


```
puppet_tags: keystone_config,keystone_domain_config
```
- d. Save and close `keystone.yaml`.
- e. Verify you see the missing tag in the `keystone.yaml` file. The following command should return '1':

```
cat /usr/share/openstack-tripleo-heat-templates/docker/services/keystone.yaml | grep 'puppet_tags: keystone_config,keystone_domain_config' | wc -l
```

2. Or, use `sed` to edit the file inline:

- a. Login to the undercloud as the stack user.
- b. Run the following command to add the missing puppet tag:


```
sed -i 's/puppet_tags\.: keystone_config/puppet_tags\.: keystone_config,keystone_domain_config/' /usr/share/openstack-tripleo-heat-templates/docker/services/keystone.yaml
```
- c. Verify you see the missing tag in the `keystone.yaml` file The following command should return '1':

```
cat /usr/share/openstack-tripleo-heat-templates/docker/services/keystone.yaml | grep 'puppet_tags: keystone_config,keystone_domain_config' | wc -l
```

BZ#1520004

It is only possible to deploy Ceph storage servers if their disk devices are homogeneous.

BZ#1525520

For deployments using OVN as the ML2 mechanism driver, only nodes with connectivity to the external networks are eligible to schedule the router gateway ports on them. However, there is a known issue that marks all nodes as eligible, which becomes a problem when the Compute nodes do not have external connectivity.

As a result, if a router gateway port is scheduled on Compute nodes without external connectivity, ingress and egress connections for the external networks will not work; in which case the router gateway port has to be rescheduled to a controller node.

As a workaround, you can provide connectivity on all your compute nodes, or you can consider deleting `NeutronBridgeMappings`, or set it to `datacentre:br-ex`. For more information, see https://bugzilla.redhat.com/show_bug.cgi?id=1525520 and https://bugzilla.redhat.com/show_bug.cgi?id=1510879.

3.3. RED HAT OPENSTACK PLATFORM 12 MAINTENANCE RELEASE 28 MARCH 2018

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

3.3.1. 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#1488855

Due to the migration from puppet-ceph to ceph-ansible for the management of Ceph using Director, old puppet hieradata (such as `ceph::profile::params::osds`) needs to be migrated to the ceph-ansible format.

Customizations for the Ceph deployment previously passed as hieradata from `*ExtraConfig` should be removed since they are ignored. Specifically, the deployment will stop if `ceph::profile::params::osds` is found to ensure the devices list has been migrated to the format expected by ceph-ansible.

Use the `CephAnsibleExtraConfig` and `CephAnsibleDisksConfig` parameters to pass arbitrary variables to ceph-ansible, such as `devices` and `dedicated_devices`.

3.3.2. Known Issues

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

BZ#1552234

There is currently a known issue where you cannot use ACLs to make a container public for anonymous access. This issue arises when sending `POST` operations to Swift that specify a `*` value in the `X-Container-Read` or `X-Container-Write` settings.

3.4. RED HAT OPENSTACK PLATFORM 12 MAINTENANCE RELEASE 20 AUGUST 2018

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

3.4.1. Enhancements

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

BZ#1502860

This update helps operators locate log files after an upgrade from a non-containerized to a containerized deployment.

If old log files are present when the upgrade begins, the old files are moved to a new file location. A `readme.txt` file is placed in the old file location. The file points to the new log file location.

For example, if a `/var/log/nova` directory exists, a `/var/log/nova/readme.txt` file is created, advising the reader to look in the `/var/log/containers/nova` directory instead.

BZ#1517278

This update prevents CPU pinning mismatches during Nova live migrations.

Prior to the update, the scheduler did not check whether the guest CPU pinning configuration was supported on the host. A mismatch of CPU pinning caused errors during bootup of the the host. This failed scenario could be repeated over a series of potential hosts.

A new condition in the `NUMATopologyFilter` filter identifies hosts with proper CPU pinning capability. If no suitable hosts are available, the migration fails quickly with an error message.

BZ#1547146

This change allows TripleO to deploy Cinder with a Dell EMC VNX backend.

BZ#1554768

Cinder volume migration between different availability zones is now supported.

BZ#1566611

Keystone user passwords generated by Heat resources such as `WaitConditionHandle` now meet more stringent regular expression-based password complexity requirements. The new passwords are 32-character random strings containing at least one uppercase and one lowercase letter, one digit, and one of the characters `!@#%&*`. These passwords should pass the standard of virtually any regular expression-based password validation.

Previously, generated passwords took the form of 32 hexadecimal digits, and thus never contained uppercase letters or special characters.

BZ#1570941

Virtual CPUs (vCPUs) can be preempted by the hypervisor kernel thread even with strong partitioning in place (`isolcpus`, `tuned`). Preemptions are not frequent, a few per second, but with 256 descriptors per virtio queue, just one preemption of the vCPU can lead to packet drop, because the 256 slots are filled during the preemption. This is the case for network functions virtualization (NFV) VMs in which the per queue packet rate is above 1 Mpps (1 million packets per second).

This release supports two new tunable options: `'rx_queue_size'` and `'tx_queue_size'`. Use these options to configure the RX queue size and TX queue size of virtio NICs, respectively, to reduce packet drop.

BZ#1571744

Nova's libvirt driver now allows the specification of granular CPU feature flags when configuring CPU models.

One benefit of this change is the alleviation of a performance degradation that has been experienced on guests running with certain Intel-based virtual CPU models after application of the "Meltdown" CVE fixes. This guest performance impact is reduced by exposing the CPU feature flag 'PCID' ("Process-Context ID") to the *guest* CPU, assuming that the PCID flag is available in the physical hardware itself.

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

3.4.2. Release Notes

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

BZ#1558148

To reduce the time spent processing security group updates in the L2 agent, `contrack` deletion is now performed in a set of worker threads instead of in the main agent thread.

BZ#1591204

A new configuration option called `bridge_mac_table_size` has been added for the neutron OVS agent. This value is set on every Open vSwitch bridge managed by the `openvswitch-neutron-agent`. The value controls the maximum number of MAC addresses that can be learned on a bridge. The default value for this new option is 50,000, which should be enough for most systems. Values outside a reasonable range (10 to 1,000,000) might be overridden by Open vSwitch.

3.4.3. Known Issues

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

BZ#1519536

You must manually discover the latest Docker image tags for current container images that are stored in Red Hat Satellite. For more information, see the Red Hat Satellite documentation: https://access.redhat.com/documentation/en-us/red_hat_satellite/6.2/html/content_management_guide/managing_container_images#managing_container_images_with_docker_tags

3.5. RED HAT OPENSTACK PLATFORM 12 MAINTENANCE RELEASE 5 DECEMBER 2018

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

3.5.1. End of Life

This release does not contain any enhancements, known issues, or release notes. Red Hat OpenStack Platform 12z4 is the final release of Red Hat OpenStack Platform 12. For more information about the OpenStack life cycle, see [Red Hat OpenStack Platform life cycle](#).

CHAPTER 4. TECHNICAL NOTES

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

4.1. RHEA-2017:3462 – RED HAT OPENSTACK PLATFORM 12.0 ENHANCEMENT ADVISORY

The bugs contained in this section are addressed by advisory RHEA-2017:3462. Further information about this advisory is available at <https://access.redhat.com/errata/RHEA-2017:3462>.

diskimage-builder

BZ#1489801

MongoDB is no longer used by Red Hat OpenStack Platform. Previously, it was used for Telemetry (which now uses Gnocchi) and Zaqr on the undercloud (which is moving to Redis). As a result, 'mongod', 'puppet-mongod', and 'v8' are no longer included.

opendaylight

BZ#1344429

This update adds the support for OpenDaylight, OVS-DPDK and OpenStack in the NetVirt/OVSDB scenario. This feature allows users to set up virtualized networks for their tenants using OpenDaylight and OVS_DPDK.

BZ#1414298

This update provides a new package of the OpenDaylight Carbon release that is used within the Red Hat OpenStack Platform 12.

BZ#1414313

With this update, the High Availability clustering is enabled for both the Neutron and the OpenDaylight controller.

BZ#1420383

This update replaces the Java based LevelDB in favour of the JNI package and provides the leveldbjni-all-1.8-15.5.el7ost.x86_64 package.

BZ#1414431

The new conntrack-based SNAT implementation, enabled by default, uses the Linux netfilter framework to do the NAT (Network Address Port Translation) and track the connection. The first packet in a traffic is passed to the netfilter to be translated with the external IP. The following packets will use the netfilter for further inbound and outbound translation. In the netfilter, the Router ID will be used as the Zone ID. Each zone tracks the connection in its own table. The rest of the

implementation remains the same.

The conntrack mode also enables the new High Availability logic that newly considers the weight associated with each switch. Also, the switch will always keep one designated NAPT port open, which improves the performance.

BZ#1450894

This update adds ping6 support to the Neutron router internal interfaces for OpenStack using OpenDaylight.

openstack-cinder

BZ#1334545

You can now set QoS IOPS limits that scale per GB size of the volume with the options "total_iops_sec_per_gb", "read_iops_sec_per_gb", and "write_iops_sec_per_gb".

For example, if you set the total_iops_sec_per_gb=1000 option, you will get 1000 IOPS for a 1GB volume, 2000 IOPS for a 2GB volume, and so on.

openstack-containers

BZ#1517903

Previously, if containers were shut down unexpectedly, Apache still left runtime files in the containers, which causes the containers to stay in a Restarting state after the host comes back up. If you use TLS everywhere, this means that the Glance and Swift services were unreachable after the host rebooted.

This fix adds runtime cleanup in the container images startup scripts. Glance and Swift services are now functioning normally after the host reboots when deployed with TLS everywhere.

openstack-neutron

BZ#1490281

Some deployments use Neutron provider bridges for internal traffic, such as traffic for AMQP, which causes bridges on boot are set to behave like normal switching. Because ARP broadcast packets use patch-ports to go between the integration bridge and the provider bridges, ARP storms to occur if more controllers were turned off ungracefully and then simultaneously booted up.

The new systemd service neutron-destroy-patch-ports now executes at the boot to remove the patch ports and break the connection between the integration bridge and the provider bridges. This prevents ARP storms, and the patch ports are then renewed after the openvswitch agent is started.

openstack-panko

BZ#1417221

The Panko service is officially deprecated in OpenStack version 12. Support for panko will be limited to usage from cloudforms only. We do not recommend using panko outside of the cloudforms use case.

openstack-tripleo-common**BZ#1276147**

This update adds support to OpenStack Bare Metal (ironic) for the Emulex hardware iSCSI (be2iscsi) ramdisk.

BZ#1434929

Previously, the `OS_IMAGE_API_VERSION` and the `OS_VOLUME_API_VERSION` environment variables were not set, which forced Glance and Cinder to fall back to the default API versions. For Cinder, this was the older v2 API.

With this update, the `overcloudrc` file now sets the environment variables to specify the API versions for Glance and Cinder.

openstack-tripleo-heat-templates**BZ#1487920**

Encrypted volumes cannot attach correctly to instances in containerized environments. The Compute service runs `"cryptsetup luksOpen"`, which waits for the udev device creation process to finish. This process does not actually finish, which causes the command to hang.

Workaround: Restart the containerized Compute service with the docker option `"--ipc=host"`.

BZ#1513109

POWER-8 (ppc64le) Compute support is now available as a technology preview.

BZ#1406102

Director now supports the creation of custom networks during the deployment and update phases. These additional networks can be used for dedicated network controllers, Ironic baremetal nodes, system management, or to create separate networks for different roles.

A single data file (`'network_data.yaml'`) manages the list of networks that will be deployed. The role definition process then assigns the networks to the required roles.

BZ#1418433

Containerized deployment of the OpenStack File Share Service (manila) is available as a technology preview in this release. By default, Manila, Cinder, and Neutron will still be deployed on bare metal machines.

■

BZ#1484467

Running Cinder services on bare metal machines and the Lscsid service in a container caused the services to have different iSCSI Qualified Name (IQN) values. Because the IQN is used to authenticate iSCSI connections, Cinder backup operations failed with an authentication error that was caused by an IQN mismatch.

With this fix, the Lscsid service now runs on bare metal, and all other services, such as containerized Nova and non-containerized Cinder, are configured to use the correct IQN.

BZ#1486995

When using an NFS back end for the Image service (glance), attempting to create an image will fail with a permission error. This is because the user ID on the host and container differ, and also because puppet cannot mount the NFS endpoint successfully on the container.

BZ#1489484

Previously, the ceph-osd package was a part of the common overcloud image, but was available only in a repository that requires the Ceph OSD entitlement. This entitlement is not required on OpenStack Controller and Compute nodes. The RPM dependency created by the ceph-osd package caused Yum update to fail when you tried to update the ceph-osd package without the ceph-osd entitlement, Yum update failed.

This fix removes the ceph-osd package from overcloud nodes that do not require the package. The ceph-osd package is now only required on Ceph storage nodes, including hyperconverged nodes that run Ceph OSD and Compute services. Yum update now succeeds on nodes that do not require the ceph-osd package. Ceph storage and hyperconverged nodes that require the ceph-osd package will still require the necessary Ceph OSD entitlement.

openstack-tripleo-puppet-elements**BZ#1270860**

Using hardcoded machine IDs in templates creates multiple nodes with identical machine IDs. This prevents the Red Hat Storage Console from identifying multiple nodes.

Workaround: Generate unique machine IDs on each node and then update the /etc/machine-id file. This will ensure that the Red Hat Storage Console can identify the nodes as unique.

BZ#1384845

When an overcloud image is shipped with 'tuned' version lower than 2.7.1-4, you should apply a manual update of the 'tuned' package to the overcloud image. If the 'tuned' version is equal to 2.7.1-4 or higher, you should provide the list of the core to 'tuned' and activate the profile, for example:

```
# echo "isolated_cores=2,4,6,8,10,12,14,18,20,22,24,26,28,30" >> /etc/tuned/cpu-partitioning-variables.conf
```

tuned-adm profile cpu-partitioning

This is a known issue until the 'tuned' packages are available in the Centos repositories.

openstack-tripleo-ui

BZ#1383576

This update adds an action to "Manage Nodes" through the director UI. This action switches nodes to a "manageable" state so the director can perform introspection through the UI.

BZ#1430885

This update increases the granularity of the deployment progress bar. This is achieved with an increase in the nesting level that retrieves the stack resources. This provides more accurate progress of a deployment.

openstack-tripleo-validations

BZ#1301549

The update adds a new validation to check the overcloud's network environment. This helps avoid any conflicts with IP addresses, VLANs, and allocation pool when deploying your overcloud.

BZ#1368512

The update adds a new validation to check the hardware resource on the undercloud before an deployment or upgrade. The validation ensures the undercloud meets the necessary disk space and memory requirements prior to a deployment or upgrade.

puppet-ironic

BZ#1489192

Previously, the DHCP server configuration file for Ironic Inspector did not handle hosts that used UEFI and iPXE, which caused some UEFI and iPXE hosts to fail to boot during Ironic Introspection. This fix updates the DHCP server file ``etc/ironic-inspector/dnsmasq.conf`` to handle UEFI and iPXE hosts, and now the hosts can properly boot during Ironic Introspection.

puppet-keystone

BZ#1404324

The token flush cron job has been modified to run hourly instead of once a day. This was changed because of issues being raised in larger deployments, as the operation would take too long and sometimes even fail because the transaction was too large. Note that this only affects deployments using the UUID token provider.

puppet-tripleo

BZ#1463355

When TLS everywhere is enabled, the HAProxy stats interface will also use TLS. As a result, you will need to access the interface through the individual node's ctlplane address, which is either the actual IP address or the FQDN (using the convention {node-name}.ctlplane.{domain}, for example, overcloud-controller-0.ctlplane.example.com). This setting can be configured by the ``CloudNameCtlplane`` parameter in ``tripleo-heat-templates``. Note that you can still use the ``haproxy_stats_certificate`` parameter from the HAProxy class, and it will take precedence if set.

BZ#1479751

Recent changes in Nova and Cinder resulted in Barbican being selected as the default encryption key manager, even when TripleO is not deploying Barbican. However, TripleO assumes that the legacy (fixed key) manager is active and selected for non-Barbican deployments. This led to broken volume encryption in non-Barbican deployments. This fix modifies the TripleO behavior to now actively configure Nova and Cinder to use the legacy key manager for non-Barbican deployments.

python-glance-store

BZ#1293435

Uploading to and downloading from Cinder volumes with Glance is now supported with the Cinder backend driver.

Note: This update does not include support for Ceph RBD. Use the Ceph backend driver to perform RBD operations on Ceph volumes.

python-openstackclient

BZ#1478287

When showing the list of Neutron security groups, the Project column referenced the tenant ID instead of the project ID. This caused the Project column to appear blank. This fix changes the behavior of the operation to get the project ID, and now the list of Neutron security groups shows the relevant project ID in the Project column.

python-os-brick

BZ#1503259

A race condition in the Python `os.path.realpath` method raised an unexpected exception. This caused an iSCSI disconnect method to unexpectedly fail. With this fix, the race condition exception is ignored. Because the symlink no longer exists, it is safe to ignore this exception. As a result, the disconnect operation succeeds, even when the race condition occurs.

python-tripleoclient

BZ#1385347

The '--controller-count' option for the 'openstack overcloud deploy' command sets the 'NeutronDhcpAgentsPerNetwork' parameter. When deploying a custom Networker role that hosts the OpenStack Networking (neutron) DHCP Agent, the 'NeutronDhcpAgentsPerNetwork' parameter might not set to the correct value. As a workaround, set the 'NeutronDhcpAgentsPerNetwork' parameter manually using an environment file. For example:

```
----  
parameter_defaults:  
  NeutronDhcpAgentsPerNetwork: 3  
----
```

This sets 'NeutronDhcpAgentsPerNetwork' to the correct value.

qemu-kvm-rhev**BZ#1498155**

Hot-unplugging Virtual Function I/O (VFIO) devices previously failed when performed after hot-unplugging a vhost network device. This update fixes the underlying code, and the VFIO device is unplugged correctly in the described circumstances.

4.2. RHEA-2018:2331 – RED HAT OPENSTACK PLATFORM 12.0 ENHANCEMENT ADVISORY AUGUST 2018

The bugs contained in this section are addressed by advisory RHEA-2018:2331. Further information about this advisory is available at <https://access.redhat.com/errata/RHSA-2018:2331.html>.

openstack-tripleo-common**BZ#1518662**

Additional non-controller upgrade attempts after a failed upgrade can fail during service validation if services are not running. To prevent such upgrade failures you can skip services validation. Pass the option "--skip-tags validation" to the Ansible invocation.

For example:

```
upgrade-non-controller.sh --upgrade compute-0 --ansible-opts "--skip-tags validation"
```

BZ#1527205

TripleO uses ceph-ansible to configure Ceph clients and servers. To reduce the undercloud memory requirement when deploying a large number of Compute nodes, the TripleO ceph-ansible fork count default was reduced from 50 to 25.

One result of the lower fork count is a reduction in the number of hosts that can be configured in parallel.

You can use a Heat environment file to override the default fork count. The following example sets the fork count to 10.

```
parameter_defaults:  
CephAnsibleEnvironmentVariables:  
DEFAULT_FORKS: '10'
```

BZ#1549139

The TripleO Derived Parameters workflow now searches for nodes in either the 'active' or 'available' states. The TripleO Derived Parameters feature searches for overcloud nodes associated with each TripleO role.

Previously, the search was limited to nodes in the 'available' state.

After the initial deployment, when nodes are typically in the 'active' state, stack updates failed because the Derived Parameters workflow did not find any nodes in the 'available' state.

BZ#1552759

The Derived Parameters workflow now supports the use of SchedulerHints to identify overcloud nodes.

Previously, the workflow could not use SchedulerHints parameters to identify overcloud nodes associated with the corresponding TripleO overcloud role. This caused the overcloud deployment to fail.

SchedulerHints support prevents these failures.

openstack-tripleo-heat-templates

BZ#1559151

Connectivity problems that occurred after OSP11-to-OSP12 upgrades have been resolved by the removal of an obsolete network configuration file.

The file was `/usr/libexec/os-apply-config/templates/etc/os-net-config/config.json`. Its presence on post-upgrade systems caused connectivity problems after a reboot on any overcloud node. Interfaces set under OVS bridges had no connectivity. For example, controller nodes were unable to rejoin the pacemaker cluster.

The upgrade process now removes the file and prevents the connectivity problems.

BZ#1559920

The file driver for Gnocchi now works as expected in containerized installations. Previously the host directory was not mounted in the container.

BZ#1571348

Database credentials are no longer logged when a transient container initializes the MySQL database on disk during a fresh overcloud deployment.

Logging verbosity was limited to prevent the logging of database credentials in the container's logs and in the journal.

BZ#1597313

A change in the libvirtd live-migration port range prevents live migration failures.

Previously, libvirtd live-migration used ports 49152 to 49215, as specified in the `qemu.conf` file. On Linux, this range is a subset of the ephemeral port range 32768 to 61000. Any port in the ephemeral range can be consumed by any other service as well.

As a result, live-migration failed with the error:
Live Migration failure: internal error: Unable to find an unused port in range 'migration' (49152-49215)

The new libvirtd live-migration range of 61152-61215 is not in the ephemeral range. The related failures no longer occur.

This completes the port change work started in BZ1573791.

BZ#1520453

An error in the `NovaSchedulerLoggingSource` variable in the `puppet/services/nova-conductor.yaml` file has been corrected to properly update logs during fluentd configuration.

Previously, `nova-scheduler.log` was tailed twice and `nova-conductor.log` was not tailed at all.

BZ#1556720

To prevent failures caused by a `gnocchi-upgrade` race condition, `gnocchi-upgrade` is now called from the bootstrap node instead of from multiple nodes.

Previously, `gnocchi-upgrade` was called from each node where `gnocchi-api` is part of the role. This sometimes resulted in failures with the error shown in the following example:
2018-03-14 12:39:39,683 [1] ERROR oslo_db.sqlalchemy.exc_filters: DBAPIError exception wrapped from (pymysql.err.InternalError) (1050, u"Table 'archive_policy' already exists")

BZ#1571435

Prior to this update, when removing the `ceph-osd` RPM from overcloud nodes that do not require the package, the corresponding Ceph OSD product key was not removed. Consequently, the `subscription-manager` would incorrectly report that the Ceph OSD product was still installed.

With this update, the script that handles removal of the `ceph-osd` RPM now also removes the Ceph OSD product key. As a result, after removing the `ceph-osd` RPM, the `subscription-manager` no longer erroneously reports the Ceph OSD product is installed.

Note: The script that removes the RPM and product key executes only during the overcloud update procedure.

BZ#1586155

OpenStack Director 13 can now successfully deploy an overcloud together with Ceph, using OpenStack 12 templates.

Prior to this update, Ceph deployment would fail during overcloud deployment step 2 because OpenStack Director failed to set the correct version of Ceph. Now OpenStack Director 12 templates always deploy the Ceph Jewel release.

BZ#1597972

This update adds the environment file `/usr/share/openstack-tripleo-heat-templates/environments/ovs-dpdk-permissions.yaml` for OVS-DPDK deployments (for new installations and minor updates).

Note: This environment file updates the parameter only for ComputeOvsDpdk role. If any other custom role is used with Ovs-DPDK then the environment file should be extended to those custom roles as well.

BZ#1502860

This update helps operators locate log files after an upgrade from a non-containerized to a containerized deployment.

If old log files are present when the upgrade begins, a `readme.txt` file is placed in the old file location. The file points to the new log file location.

For example, if a `/var/log/nova` directory exists, a `/var/log/nova/readme.txt` file is created, advising the reader to look in the `/var/log/containers/nova` directory instead.

BZ#1508867

This update adds the service `OS::TripleO::Services::NovaMigrationTarget` to the service list of the ComputeOvsDpdk role in the `roles_data.yaml`. Prior to this update, the omission of the service caused Nova live migration to fail on the ComputeOvsDpdk roles.

Before starting a minor update, ensure the service is present in the ComputeOvsDpdk role of the `roles_data.yaml` file.

BZ#1547146

This change allows TripleO to deploy Cinder with a Dell EMC VNX backend.

BZ#1585362

The TripleO environment files used for deploying Cinder's Netapp backend have been updated in this release to allow successful deployment of a Cinder Netapp backend.

Prior to this update, obsolete data caused the overcloud deployment to fail.

BZ#1589951

The default age for purging deleted database records has been corrected so that deleted records

are purged from Cinder's database.

Previously, the CinderCronDbPurgeAge value for Cinder's purge cron job used the wrong value and deleted records were not purged from Cinder's database when they reached the desired default age.

BZ#1573808

To enable the neutron-lbaas dashboard:

1. Enable the dashboard in the Horizon configuration file in all controller nodes:

File: `/var/lib/config-data/puppet-generated/horizon/etc/openstack-dashboard/local_settings`

```
OPENSTACK_NEUTRON_NETWORK = {  
'enable_distributed_router': False,  
'enable_firewall': False,  
'enable_ha_router': False,  
'enable_lb': True, <-----
```

2. Restart the horizon container:

```
# docker restart horizon
```

A new "Load Balancers" tab will appear under the "Network" menu. The URL is `http://<controller-vip>/dashboard/project/ngloadbalancersv2`

puppet-nova

BZ#1571744

Nova's libvirt driver now allows the specification of granular CPU feature flags when configuring CPU models.

One benefit of this change is the alleviation of a performance degradation that has been experienced on guests running with certain Intel-based virtual CPU models after application of the "Meltdown" CVE fixes. This guest performance impact is reduced by exposing the CPU feature flag 'PCID' ("Process-Context ID") to the *guest* CPU, assuming that the PCID flag is available in the physical hardware itself.

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

puppet-tripleo

BZ#1528632

Prior to this update, running a "stack update" operation on an existing stack to reassess the state of Heat resources caused a failure in container `docker-puppet-rabbitmq`. This failure prevented users from running stack update operations.

This update fixes the issue by changing the way puppet configuration is done in the `rabbitmq` container `docker-puppet-rabbitmq`.

BZ#1585149

This update allows a non-containerized OpenStack service to connect to the Ceph cluster.

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

Puppet now sets the file ACLs mask for the CephX keyrings so that it is allowed to grant read permissions to specific users.

BZ#1533511

This fix prevents a potential failure of ceilometer-upgrade during an OpenStack upgrade.

Prior to this fix, the ceilometer-upgrade sometimes failed during an OSP11-OSP12 upgrade because it ran before gnocchi-upgrade.

If you upgraded to OSP12 without this fix and ceilometer-upgrade failed, delete the `/etc/gnocchi/gnocchi.conf` file from the bootstrap node and re-run the upgrade process with the fixed package.

BZ#1590953

This update fixes an issue that prevented users from configuring Netapp NFS mount options via the TripleO Heat parameter.

Prior to this update, the Cinder Netapp backend ignored the `CinderNetappNfsMountOptions` TripleO Heat parameter, preventing configuration of the Netapp NFS mount options via the TripleO Heat parameter.

The code responsible for handling Cinder's Netapp configuration no longer ignores the `CinderNetappNfsMountOptions` parameter. The `CinderNetappNfsMountOptions` parameter correctly configures Cinder's Netapp NFS mount options.

BZ#1599410

During a version upgrade, Cinder's database synchronization is now executed only on the bootstrap node. This prevents database synchronization and upgrade failures that occurred when database synchronization was executed on all Controller nodes.

python-os-brick**BZ#1572572**

OS-Brick FC host bus adapter (HBA) scans have been limited to prevent the addition of unwanted devices.

Previously, the OS-Brick FC code always scanned all present HBAs.

Now the following limits apply:

- If an initiator map is present, only the mapped HBAs are scanned
- In the case of a single WWNN for all ports, only the connected HBAs are scanned
- Else, all HBAs are scanned with wildcards

4.3. RHEA-2018:2332 – RED HAT OPENSTACK PLATFORM 12.0 SECURITY ADVISORY AUGUST 2018

The bugs contained in this section are addressed by advisory RHSA-2018:2332. Further information about this advisory is available at <https://access.redhat.com/errata/RHSA-2018:2332.html>.

openstack-nova

BZ#1570941

Virtual CPUs (vCPUs) can be preempted by the hypervisor kernel thread even with strong partitioning in place (isolcpus, tuned). Preemptions are not frequent, a few per second, but with 256 descriptors per virtio queue, just one preemption of the vCPU can lead to packet drop, because the 256 slots are filled during the preemption. This is the case for network functions virtualization (NFV) VMs in which the per queue packet rate is above 1 Mpps (1 million packets per second).

This release supports two new tunable options: 'rx_queue_size' and 'tx_queue_size'. Use these options to configure the RX queue size and TX queue size of virtio NICs, respectively, to reduce packet drop.

BZ#1558706

Previously, the ability to set an admin password to the metadata service was not implemented for the libvirt driver causing the 'nova get-password' command to return nothing.

This release enables setting an admin password to the metadata service for the libvirt driver. The admin password is saved to the metadata service, and the 'nova get-password' command returns that password.

BZ#1563109

This update slows the initial stages of live migrations to eliminate packet loss. Previously, instances with LinuxBridge VIFs experienced packet loss during live migration. Neutron did not have enough time to complete the plugging of the VIFs and related networking infrastructure on the destination during live migration.

Live migrations are now initially slowed to ensure Neutron has adequate time to wire up the VIFs on the destination. Once complete, Neutron sends an event to Nova, returning the migration to full speed. This requires Neutron 11.0.4 or greater on Pike when used with LinuxBridge VIFs to pick up the Icb039ae2d465e3822ab07ae4f9bc405c1362afba bugfix.

BZ#1579785

Prior to this update, to re-discover a compute node record after deleting a host mapping from the API database, the compute node record had to be manually marked as unmapped. Otherwise, a compute node with the same hostname could not be mapped back to the cell from which it was

removed.

With this update, the compute node record is automatically marked as unmapped when you delete a host from a cell, enabling a compute node with the same hostname to be added to the cell during host discovery.

BZ#1517278

This update prevents CPU pinning mismatches during Nova live migrations.

Prior to the update, the scheduler did not check whether the guest CPU pinning configuration was supported on the host. A mismatch of CPU pinning caused errors during bootup of the the host. This failed scenario could be repeated over a series of potential hosts.

A new condition in the NUMATopologyFilter filter identifies hosts with proper CPU pinning capability. If no suitable hosts are available, the migration fails quickly with an error message.

BZ#1539703

This update prevents an unintended bypass of the schedule filters This update prevents an unintended bypass of the schedule filters that could occur after the scheduler refused a rebuild request sent by nova.

If a user rebuilds an instance with a new image, the change from old image to new image causes nova to send the rebuild request to the scheduler to make sure it is allowed according to the scheduler filters.

Prior to this update, if the scheduler refused the request, the instance's image reference was not rolled back to the original image. This caused an inconsistency between the original image actually in use by the instance, and the new image reference saved in the database. As a result, a second rebuild request with the same new image would bypass the scheduler and be allowed, because the image in the rebuild request was the same as the instance's image in the database, even though the real image in use by the instance was the old original image. This bypass of scheduler filters was considered a security flaw.

As of this update, when a rebuild request is refused by the scheduler, the image reference is rolled back to the original. If another rebuild request is made with the same new image, it is correctly identified as being different from the instance's current image and the request is send to the scheduler.

BZ#1547578

Prior to this update, a volume detach operation performed under certain failure scenarios could result in the removal of a volume's libvirt definition without full removal of the associated logical volume (LUN) from the host. This allowed Cinder to incorrectly perform subsequent operations while the compute host still had active paths to the device.

As of this update, even under a failure scenario, Nova compute attempts to disconnect the LUN from the host. The result is a better release of the logical volume on the host.

