



Red Hat OpenStack Platform 14

Release Notes

Release details for Red Hat OpenStack Platform 14

Red Hat OpenStack Platform 14 Release Notes

Release details for Red Hat OpenStack Platform 14

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

Legal Notice

Copyright © 2019 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

CHAPTER 1. INTRODUCTION	3
1.1. ABOUT THIS RELEASE	3
1.2. REQUIREMENTS	3
1.3. DEPLOYMENT LIMITS	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 OPERATING SYSTEMS	4
1.8. HYPERVISOR SUPPORT	4
1.9. CONTENT DELIVERY NETWORK (CDN) REPOSITORIES	4
1.10. PRODUCT SUPPORT	6
CHAPTER 2. TOP NEW FEATURES	7
2.1. RED HAT OPENSTACK PLATFORM DIRECTOR	7
2.2. BARE METAL SERVICE	8
2.3. CEPH STORAGE	8
2.4. COMPUTE	8
2.5. METRICS AND MONITORING	9
2.6. NETWORK FUNCTIONS VIRTUALIZATION	9
2.7. OPENSTACK NETWORKING	10
2.8. STORAGE	10
2.9. TECHNOLOGY PREVIEWS	11
2.9.1. New Technology Previews	11
CHAPTER 3. RELEASE INFORMATION	14
3.1. RED HAT OPENSTACK PLATFORM 14 GA	14
3.1.1. Enhancements	14
3.1.2. Technology Preview	18
3.1.3. Release Notes	19
3.1.4. Known Issues	20
3.1.5. Deprecated Functionality	24
3.2. RED HAT OPENSTACK PLATFORM 14 MAINTENANCE RELEASE 13 MARCH 2019	24
3.2.1. Enhancements	25
3.2.2. Known Issues	25
3.3. RED HAT OPENSTACK PLATFORM 14 MAINTENANCE RELEASE 30 APRIL 2019	26
3.3.1. Enhancements	26
3.3.2. Known Issues	26
3.3.3. Deprecated Functionality	27
CHAPTER 4. TECHNICAL NOTES	28
4.1. RHEA-2019:0045 — OPENSTACK DIRECTOR BUG FIX ADVISORY	28

CHAPTER 1. INTRODUCTION

1.1. ABOUT THIS RELEASE

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

<https://releases.openstack.org/rocky/index.html>.

Red Hat OpenStack Platform uses components from other Red Hat products. See the following link(s) 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 link for more details on the add-on: <http://www.redhat.com/products/enterprise-linux-add-ons/high-availability/>. See the following link for details on the package versions to use in combination with Red Hat OpenStack Platform: <https://access.redhat.com/site/solutions/509783>

1.2. REQUIREMENTS

This version of Red Hat OpenStack Platform runs on the most recent fully supported release of Red Hat Enterprise Linux.

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 runs on 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 OPERATING SYSTEMS

For a list of the 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

This release of the Red Hat OpenStack Platform is supported only 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. This release of the Red Hat OpenStack Platform runs with Ironic.

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

1.9. CONTENT DELIVERY NETWORK (CDN) REPOSITORIES

This section describes the repository settings required to deploy Red Hat OpenStack Platform 14.

You can install Red Hat OpenStack Platform 14 through the Content Delivery Network (CDN). To do so, configure **subscription-manager** to use the correct repositories.

Run the following command to enable a CDN repository:

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

Run the following command to disable a CDN repository:

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

Table 1.1. Required Repositories (x86_64)

Repository Name	Repository Label
Red Hat Enterprise Linux 7 Server (RPMS)	<code>rhel-7-server-rpms</code>
Red Hat Enterprise Linux 7 Server - RH Common (RPMS)	<code>rhel-7-server-rh-common-rpms</code>
Red Hat Enterprise Linux High Availability (for RHEL 7 Server)	<code>rhel-ha-for-rhel-7-server-rpms</code>
Red Hat OpenStack Platform 14 for RHEL 7 (RPMS)	<code>rhel-7-server-openstack-14-rpms</code>
Red Hat Enterprise Linux 7 Server - Extras (RPMS)	<code>rhel-7-server-extras-rpms</code>

Table 1.2. Optional Repositories (x86_64)

Repository Name	Repository Label
Red Hat Enterprise Linux 7 Server - Optional	<code>rhel-7-server-optional-rpms</code>
Red Hat OpenStack Platform 14 Operational Tools for RHEL 7 (RPMS)	<code>rhel-7-server-openstack-14-optools-rpms</code>

Table 1.3. Required Repositories (ppc64le)

Repository Name	Repository Label
Red Hat Enterprise Linux for IBM Power, little endian	<code>rhel-7-for-power-le-rpms</code>
Red Hat OpenStack Platform 14 for RHEL 7 (RPMS)	<code>rhel-7-server-openstack-14-for-power-le-rpms</code>

Repositories to Disable

The following table outlines the repositories you must disable to ensure Red Hat OpenStack Platform 14 functions correctly.

Table 1.4. Repositories to Disable

Repository Name	Repository Label
Red Hat CloudForms Management Engine	<code>"cf-me-*</code>
Red Hat Enterprise Virtualization	<code>"rhel-7-server-rhev*"</code>

Repository Name	Repository Label
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:

- Product documentation.
- Knowledge base articles and solutions.
- Technical briefs.
- Support case management.

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

Mailing Lists

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

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

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

CHAPTER 2. TOP NEW FEATURES

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

2.1. RED HAT OPENSTACK PLATFORM DIRECTOR

This section outlines the top new features for the director.

Ansible-driven deployment using director

With this release, Ansible is now integrated into the deployment process. This allows you to use Ansible tooling for certain tasks, including dry-run, targeted tasks execution, among others:

- Ansible now performs software configuration in Director, using the feature name **config-download**.
 - This capability was previously Tech Preview in Red Hat OpenStack Platform 13, and is now entering General Availability in Red Hat OpenStack Platform 14.
- Heat still defines the software configuration, but does not apply it:
 - The configuration is made available by Heat.
 - **ansible-playbook** downloads the configuration and then applies it. The undercloud serves as the Ansible control node.

Advanced subscription manager

You can now define which roles will consume a particular subscription/pool. This means that you can use only the subscriptions you need.

- New Ansible role added to manage subscriptions.
- Richer management options.
- Ability to assign subscriptions/pools per role.

Removal of Ceph and OpenStack services from Overcloud images

As a result of the container implementation, these services have been changed:

- Removal of OpenStack services.
- Removal of Ceph packages.
- OpenStack clients are still installed.
- Minimal OpenStack content required for deployment.
 - Note that **python-heat-agents** are still installed.
- Ceph entitlements are no longer needed for all nodes (an alternative product SKU is available).

Automated container image building

You can use director to build a customized container image based on your own definition, allowing you to avoid extra manual steps before deployment.

More Ansible roles to automate image customization

- New Ansible role to automate image customization.
- The operator defines the docker file.
- Director can build an extended container image based on a given definition, and push it to the registry.

Containerized and unified undercloud

This release uses a unified installation procedure for the undercloud and overcloud, letting you take advantage of overcloud deployment features.

- No need to learn or maintain separate procedures.
- The undercloud runs in containers.
- Improvements have been added to the overcloud deploy process.
- You can define the required set of services.
- You may find that this approach makes it easier to evaluate Red Hat OpenStack Platform.

2.2. BARE METAL SERVICE

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

Bare metal deployment options

Director in OSP 14 can deploy the OpenShift Container Platform on bare metal nodes on RHEL using the openshift-ansible templates under the hood, transparently for the operator, who has to interact only with director. Director will also allow to add and remove OCP nodes accordingly.

2.3. CEPH STORAGE

This section outlines the top new features for Ceph Storage.

Create and manage a multi-tier Ceph storage via director

Using OpenStack director, you can deploy different Red Hat Ceph Storage performance tiers by adding new Ceph nodes dedicated to a specific tier in a Ceph cluster.

For example, you can add new object storage daemon (OSD) nodes with SSD drives to an existing Ceph cluster to create a Block Storage (cinder) back end exclusively for storing data on these nodes. A user creating a new Block Storage volume can then choose the desired performance tier: either HDDs or the new SSDs.

This type of deployment requires Red Hat OpenStack Platform director to pass a customized CRUSH map to ceph-ansible. The CRUSH map allows you to split OSD nodes based on disk performance, but you can also use this feature for mapping physical infrastructure layout.

Improved integration with ceph-ansible

This release rewrites director's ceph-ansible integration to work with the new config-download feature to provide a better user experience. Users can more easily troubleshoot director Ceph deployments by using the Ansible `external_deploy_steps` tag.

2.4. COMPUTE

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

TX/RX Queue Sizing

You can configure the queue size of TX and RX traffic for **libvirt** and **virtio** interfaces. You can define the queue size for each host or each guest as needed, to improve performance and handle increased traffic use-cases. The parameter for the TX/RX queue size is available in the relevant role data file before deployment, and in the **nova.conf** file after deployment.

Trusted Virtual Functions (VFs) for SR-IOV

You can designate instances as trusted, which then enables you to change the MAC address of the VF and enable promiscuous mode directly from the guest instance. These functions help you configure failover VFs for instances directly from the instance.

NFS backend for Nova

You can mount Compute instances from an NFS export, and maintain a shared NFS storage backend for instances. This functionality works in a similar way to the NFS storage backend for Glance and Cinder.

Reserved huge pages

You can allocate huge pages to specific Compute nodes to support high-performance workloads. To reserve huge pages for specific nodes, set the **reserved_huge_pages** parameter in the Director before deployment. The configuration is then available in the **nova.conf** file after deployment.

2.5. METRICS AND MONITORING

This section outlines the top new features and changes for the metrics and monitoring components.

2.6. NETWORK FUNCTIONS VIRTUALIZATION

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

Configure emulator threading per host

You can configure deterministic performance by not over-committing a vCPU in QEMU, in order to avoid spurious packet drops. In a given OSP-d composable role, you can now choose which host CPUs will run the QEMU emulator threads. For example:

```
parameter_defaults:
  ComputeOvsDpdkParameters:
    NovaComputeCpuSharedSet: "0-1"
```

Red Hat's recommendation is to use the same CPU set as the host (non-isolated CPUs):

```
HostCpusList: "0-1"
```

This is then activated per VM flavor:

```
hw:emulator_threads_policy=share
```

Use introspection to calculate NFV parameters

You can use introspection to calculate certain SR-IOV and OVS-DPDK director parameters. This is expected to ease the deployment of NFVi. For example:

```
workflow_parameters: tripleo.derive_params.v1.derive_parameters:  
  num_phy_cores_per_numa_node_for_pmd: 2  
  huge_page_allocation_percentage: 90
```

2.7. OPENSTACK NETWORKING

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

ML2/OVS to ML2/OVN Migration

This update provides an in-place migration strategy from ML2/OVS to ML2/OVN in either ovs-firewall or ovs-hybrid mode for an OpenStack deployment with director.

Neutron internal DNS resolution

The DHCP agent now passes `dns_domain` to the network's `dnsmasq` process, in turn passing it to the instances.

OVN services status report

The `openstack network agent list` command now reports on all OVN services and their status.

Octavia (LBaaS) improved deployment

The latest Octavia images are automatically pushed during update or upgrade.

Octavia controller container health monitoring

This release introduces the ability to monitor Octavia container service VM health.

Multi-tenant BMaaS with new Ansible Networking ML2 plugin

This release allows multiple tenants to use nodes in an isolated fashion.

2.8. STORAGE

Block Storage - Support for signed Glance images

The Block Storage Service (cinder) automatically validates the signature of any downloaded, signed image during volume creation. The signature is validated before the image is written to the volume. Users now have stronger assurances of the integrity of the image data they are using to create volumes. This feature does not work with Ceph storage.

Block Storage - Migration between cinder availability zones

Volume migration across availability zones was added to the Block Storage service (cinder) so users can migrate volume from one availability zone to another.

Block Storage - Cinder backup NFS support

Prior to this release, the Red Hat OpenStack Platform director could only deploy the Object Storage service (swift) or Red Hat Ceph Storage as a backup back end. The Block Storage service (cinder) backup NFS profile support was introduced in director to expand Red Hat OpenStack Platform deployment to support Ceph, NFS, and Swift as backup targets. Now, director can deploy NFS as the back end for the backup service using the `CinderBackupBackEnd` parameter in the `cinder-backup.yaml` Heat template.

Block Storage - Optimized RBD to RBD migration

This release implements an optimized Ceph RBD to RBD block-level volume migration to take advantage of the underlying Ceph back end capabilities when both back ends (source and target) reside on the same Ceph cluster. This feature enables faster and more efficient data migration

operations, such as when you retire old hardware, move between tiers, and so forth.

Data Processing - S3 compatible object stores

This release introduces Hadoop support for S3-compatible object stores in the Data Processing service (sahara). This feature follows on the efforts to make data sources and job binaries "pluggable". The S3 support is an additional alternative to the existing HDFS, swift, MapR-FS, and manila storage options.

Image Service - Transparent image conversion

When importing a new image, the Image service (glance) now automatically converts the image format from QCOW2 to RAW as the destination format (without intervention) when using Ceph as the backend for the Image service.

Object Storage - Object Storage S3 API by default

The S3 API is regarded by the industry as the defacto object storage standard API. Red Hat Openstack Object Storage service (swift) supported the S3 API by using the Swift3 middleware, as a post-deployment manual operation. Starting with this release, the Swift3 middleware is set by default on an overcloud deployment.

Shared File System - Manila share-type quotas support

Cloud administrators can now define the quota for the number of shares for a given share type. This functionality is similar to the one offered by the Block Storage service (cinder) for quotas per volume type. In setups with multiple share types, the per share type quota allows resource providers to have better control over the provisioned resources.

Shared File System - User message support

Until this release, if manila operations failed asynchronously (e.g., to create share or create share group), the user did not receive any detailed information. This new capability provides more information to users about failed asynchronous operations to better troubleshoot their errors and possibly recover, without cloud administrator intervention.

2.9. TECHNOLOGY PREVIEWS

This section outlines features that are in technology preview in Red Hat OpenStack Platform 14.



NOTE

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

2.9.1. New Technology Previews

The following new features are provided as technology previews:

Virtual GPU (vGPU) support for instances

To access GPU-based rendering on your guest instances, you can define and manage virtual GPU (vGPU) resources according to your available physical GPU devices. This configuration allows you to more effectively divide the rendering workloads between all your physical GPU devices, and to have more control over scheduling, tuning, and monitoring your vGPU-enabled guest instances.

**NOTE**

- Currently vGPU support is provided as a technical preview only for NVIDIA GRID vGPU devices. You must comply with the NVIDIA GRID licensing requirements.
- Only one vGPU type is supported per physical GPU, and only one vGPU resource is supported per guest instance.

NUMA-aware vSwitches

OpenStack Compute now takes into account the NUMA node location of physical NICs when launching a Compute instance. This helps to reduce latency and improve performance when managing DPDK-enabled interfaces.

OpenDaylight - VXLAN DSCP inheritance

OpenDaylight supports DSCP inheritance, whereby DSCP markings on the inner IP header are replicated to the DSCP markings on the outer IP header for VXLAN encapsulated packets. With this feature, tenant traffic is forwarded over VXLAN tunnels based on DSCP markings from the tenant.

Automatic restart of instances on Compute node reboot

You can now configure automatic restart of instances on a Compute node even if you do not migrate the instances first. The Compute service and the libvirt-guests agent can be configured to gracefully shut down the instances and then start the instances again after the Compute node reboots.

The following parameters are available:

- **NovaResumeGuestsStateOnHostBoot** (True/False)
- **NovaResumeGuestsShutdownTimeout** (default 300s)

Skydive - network visualization suite

Skydive is a complete network visualization and monitoring suite, targeted for the cloud operator. Features include the following.

- Network topology discovery
 - Live and historical analysis
 - Metrics and alerting system
 - Packet generator for tracing and validating network infrastructure
- Skydive is fully integrated with OSP director. It supports all OVS based systems, including OVN and OpenDaylight. It exposes REST API, command line interface (CLI), and Web UI.

Metrics and Monitoring - Service Assurance Framework

This releases adds a Technology Preview of the Service Assurance Framework, allowing for metrics and events monitoring at scale. This is a platform-based approach to Metrics and Monitoring, and is based on the following elements:

- Collectd plug-ins for infrastructure and OpenStack service monitoring.
- AMQ Interconnect direct routing (QDR) message bus.
- Prometheus Operator database/management cluster.
- Ceilometer/Gnocchi for chargeback/capacity planning only.

Block Storage - Attach a volume to multiple hosts

This release adds the ability to attach a volume to multiple hosts or servers simultaneously in both cinder and nova with read/write (RW) mode when this feature is supported by the back end driver. This feature addresses the clustered application workloads use case that typically requires active/active or active/standby scenarios.

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 14 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#1241017

This update adds hostname and network name to the output of the 'openstack port list' command.

The additional information makes it easier to associate Neutron port and IP addresses with a particular host.

BZ#1402584

With this update, the libvirt compute driver now allows users to create instances with trusted SR-IOV virtual functions. When trusted, a VF can perform certain operations, such as modifying the VF's MAC address in the guest.

Interface bonding requires that all slaves use the same MAC address, which in turn requires MAC address modifications on one of the VFs during a failover. Because MAC address altering is a privileged operation, participating VFs must be trusted in order to successfully configure bonding in the guest.

Administrators can now configure trusted mode for VFs. This requires two steps. First, the 'trusted' value of the '[pci] passthrough_whitelist' JSON configuration option in nova.conf must be set to 'true'. For example:

```
[pci]
passthrough_whitelist = {"devname": "eth0", "trusted": "true",
                        "physical_network": "sriovnet1"}
```

Then, when creating the port, 'trusted=true' must be set for the binding profile. For example:

```
$ neutron port-create <net-id> \
  --name sriov_port \
  --vnic-type direct \
  --binding:profile type=dict trusted=true
```

Because trusted mode only applies to SR-IOV VFs, the 'vnic-type' must be one of 'hw_veb' or 'direct'.

BZ#1410195

Heat templates now include the `CephClusterName` parameter. This parameter enables you to customize the Ceph cluster name, which you might need to do if you use an external Ceph cluster or a Ceph RBDMirror.

BZ#1462048

With this update, Users can create application credentials to allow their applications to authenticate to keystone.

See

https://docs.openstack.org/keystone/latest/user/application_credentials.html

BZ#1469073

Feature:

Add the CPUWeigher weigher for nova-scheduler.

Reason:

The CPUWeigher allows operators to configure a stacking or spreading policy for vCPUs.

Result:

Operators can enable the CPUWeigher and configure a stacking (use all vCPUs on one node first) or spreading (attempt to use a roughly equal amount of vCPUs from all hosts) policy.

BZ#1512941

This update supports two new tunable options that can be used to reduce packet drop.

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 update 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#1521176

Instance resources have three new attributes: `launched_at`, `deleted_at`, `created_at` to track the exact time that Nova creates/launches/deletes instances.

BZ#1523328

OpenStack director now uses Ansible for software configuration of the overcloud nodes. Ansible provides a more familiar and debuggable operator experience during overcloud deployment. Ansible is used to replace the communication and transport of the software configuration deployment data between heat and the heat agent (`os-collect-config`) on the overcloud nodes.

Instead of `os-collect-config` running on each overcloud node and polling for deployment data from heat, the Ansible control node applies the configuration by running an `ansible-playbook` with an Ansible inventory file and a set of playbooks and tasks. The Ansible control node (the node running `ansible-playbook`) is the undercloud by default.

BZ#1547708

OpenStack Sahara now supports Cloudera Distribution Hadoop (CDH) plugin 5.13.

BZ#1547710

This update adds support of s3-compatible object stores for OpenStack Sahara.

BZ#1547954

With this release, 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 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 ```nova.conf``` for usage details.

BZ#1562171

This update introduces multi-tenant bare metal networking with the "neutron" network interface.

By configuring the bare metal nodes with the "neutron" network interface, an operator can enable the users to use isolated VLAN networks for provisioning and tenant traffic on bare metal nodes.

BZ#1639759

A new tripleo heat template has been added to support QDR. Enabling the metrics_qdr service deploys QDR service on all overcloud nodes, which will be used for routing metrics data from collectd service running on each overcloud node.

The relevant heat template parameters are shown below.

```

MetricsQdrPort:
  default: '5666'
  description: Service name or port number on which the qdrouterd will
  accept
                connections. This argument must be string, even if the
numeric
                form is used.
  type: string
MetricsQdrUsername:
  default: 'guest'
  description: Username which should be used to authenticate to the
  deployed
                qdrouterd.
  type: string
MetricsQdrPassword:
  default: 'guest'
  description: Password which should be used to authenticate to the
  deployed
                qdrouterd.
  type: string
  hidden: true
MetricsQdrConnectors:
  default: []
  description: Connectors configuration (array of hashes).
  type: json
MetricsQdrAddresses:
  default:
    - prefix: 'collectd/notify'
      distribution: multicast
    - prefix: 'collectd/telemetry'
      distribution: multicast
  description: Addresses configuration (array of hashes).
  type: json
MetricsQdrUseSSL:
  default: false
  description: Set to true if it is required to use SSL or TLS on
                the connection for listener.
  type: boolean
MetricsQdrUseEncryption:
  default: false
  description: Set to true if it is required to encrypt connection to the
  peer
                for listener.
  type: boolean
MetricsQdrSaslMechanisms:
  default: 'ANONYMOUS'
  description: List of accepted SASL auth mechanisms for listener in
  format
                of comma separated list.

```

```

    type: string
MetricsQdrSslCertDb:
    default: ''
    description: Path to SSL certificate db for listener.
    type: string
MetricsQdrSslCertFile:
    default: ''
    description: Path to SSL certificate file for listener.
    type: string
MetricsQdrSslKeyFile:
    default: ''
    description: Path to SSL private key file for listener.
    type: string
MetricsQdrSslPwFile:
    default: ''
    description: Path to SSL password file for certificate key for
listener.
    type: string
MetricsQdrSslPassword:
    default: ''
    description: SSL password to be supplied for listener.
    type: string
MetricsQdrTrustedCerts:
    default: ''
    description: Path to file containing trusted certificates for listener.
    type: string

```

[BZ#1654123](#)

Red Hat OpenStack Platform 14 is now supported on IBM POWER9 CPUs. This support is provided with the `rhosp-director-images-ppc64lep9` and `rhosp-director-images-ipa-ppc64lep9` packages.

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#1033180](#)

This release adds a Technology Preview of the ability to attach a volume to multiple hosts or servers simultaneously in both cinder and nova with read/write (RW) mode when this feature is supported by the back end driver. This feature addresses the clustered application workloads use case that typically requires active/active or active/standby scenarios.

[BZ#1550668](#)

This feature enables forwarding tenant traffic based on DSCP marking from tenants encapsulated in the VXLAN IP header. This feature is a technology preview for OSP14.

[BZ#1614282](#)

You can now configure automatic restart of instances on a Compute node if the compute node reboots without first migrating the instances. Nova and the libvirt-guests agent can be configured to gracefully shut down the instances and start them when the Compute node reboots.

New parameters:

NovaResumeGuestsStateOnHostBoot (True/False)

NovaResumeGuestsShutdownTimeout (default 300s)

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#1601613

The default value of `--http-boot` changed from `/httpboot` to `/var/lib/ironic/httpboot` as containerized Ironic services expect.

BZ#1614810

With this update, logrotate's copytruncate is used by default for containerized services logs rotation. The default period to keep old logs remains unchanged (14 days).

BZ#1640095

OpenStack Rally, previously included as a technical preview, is removed from this release.

BZ#1649679

When you use the web-download feature, the staging area - defined in the configuration using the `node_staging_uri` option - is not cleaned up properly. Ensure that `file` is part of the `stores` configuration option in the `glance_store` section of the `glance-api.conf` file.

BZ#1654405

When you use the image conversion feature, ensure that `file` is part of the `stores` configuration option in the `glance_store` section of the `glance-api.conf` file.

BZ#1654408

For glance image conversion, the glance-direct method is not enabled by default. To enable this feature, set `enabled_import_methods` to `[glance-direct,web-download]` or `[glance-direct]` in the DEFAULT section of `glance-api.conf`.

BZ#1654413

Glance image conversion is not enabled by default on a new install of Red Hat OpenStack Platform 14. To use this feature, edit the `glance-image-import.conf` file.

In the `image_import_opts` section, insert the following line:
`image_import_plugins = ['image_conversion']`

[BZ#1662042](#)

OpenDaylight does not support IPv6 for tenant or provider networks. Therefore, use only IPv4 networks. You may experience issues related to floating IPs if IPv6 networks are used along with IPv4 networks.

3.1.4. Known Issues

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

[BZ#1516911](#)

The `OvsDpdkMemoryChannels` parameter cannot be derived through the DPDK derive parameters workflow. The value is set to 4 by default. You can change that value in your custom environments file to match your hardware.

[BZ#1579052](#)

When Octavia is configured to use a small Nova flavor, Amphorae (Nova instances) are created successfully but load balancers can get stuck in PENDING state for about 25 minutes. Instead, the load balancer should go to error state and the Amphorae should be deleted.

As a workaround for small Nova flavors, tune Octavia configurations `"connection_max_retries"`, `"connection_retry_interval"`, `"build_active_retries"` and `"build_retry_interval"` in section `[haproxy_amphora]` to a more reasonable production values. This will cause load balancers will transition from PENDING to ERROR state faster with a small Nova flavor.

[BZ#1630480](#)

Workflow triggers for generating Openstack rc files are hardcoded in `python-tripleoclient`.

As a result, OpenStack-specific workflows are triggered after director deploys OpenShift. Users can see OpenStack-specific URLs in stdout and OpenStack rc files created.

[BZ#1639495](#)

There is currently a known issue with fernet token rotation where the keys are not automatically deployed onto the overcloud. The workflow task ``tripleo.fernet_keys.v1.rotate_fernet_keys`` generates the keys but they are not successfully pushed to the overcloud. This issue is expected to be addressed in a future release. If you plan to perform rotation before this update, you can choose to follow one of these workarounds:

* Start `os-collect-config` on the overcloud nodes before running the

rotation. You can then stop it afterwards if you do not need it for anything else.

* Enable `os-collect-config` on all overcloud nodes. You can choose to disable it once the update with the fix is released.

NOTE: If you do not need to rotate keys before the update comes out, then you do not need to do anything.

BZ#1640021

Deploying the overcloud with a Gnocchi file backend might fail due to access permissions to the `/var/lib/gnocchi/` directory.

Workaround: Before you deploy the overcloud, set the permissions to the directory in the `openstack/tripleo-heat-templates/docker/services/gnocchi-api.yaml` file as follows:

```
- path:
  list_join:
    - "/"
    - - {get_param: GnocchiFileBasePath}
      - "tmp"
  owner: gnocchi:gnocchi
  perm: '0600'
  recurse: true
```

BZ#1640382

On a director-deployed OpenShift environment, the GlusterFS playbooks auto-generate a new heketi secret key for each run. As a result of this, operations such as scale out or configuration changes on CNS deployments fail.

As a workaround, complete the following steps:

1. Post-deployment, retrieve the heketi secret key. Use this command on one of the master nodes:

```
sudo oc get secret heketi-storage-admin-secret --namespace glusterfs -o json | jq -r .data.key | base64 -d
```

2. In an environment file, set the following parameters to that value:

```
openshift_storage_glusterfs_heketi_admin_key
openshift_storage_glusterfs_registry_heketi_admin_key
```

As a result of this workaround, operations such as scale out or configuration changes on CNS deployments work as long as the parameters were manually extracted.

BZ#1640804

When you restart all three controller nodes, it might not be possible to launch tenant instances in the overcloud. A "DuplicateMessageError" message is logged in the overcloud logs.

As a workaround, on one of the overcloud controllers, run this command:
`pcs resource restart rabbitmq-bundle`

BZ#1643657

■

For proxying requests to the routers on Infra nodes, director sets up port 443 on the HAProxy instance running on master nodes. Port 443 cannot be used on OpenShift master nodes for binding the OpenShift API. OpenShift API cannot be configured on port 443 on a director deployed OpenShift environment.

BZ#1644889

The overcloud-full image provided by director causes an RPM conflict with the python-setuptools provided by the OpenShift repos. Any post-deployment yum update on the OpenShift nodes fails with a broken dependency issue.

To fix this, run the following on the undercloud:

```
source ~/stackrc
tripleo-ansible-inventory --stack openshift --static-yaml-inventory
/home/stack/openshift_inventory.yaml
export ANSIBLE_HOST_KEY_CHECKING=False
ansible -i openshift_inventory.yaml -m shell -b -a 'rpm -e --nodeps
python-setuptools-0.9.8-7.el7.noarch; yum -y install python-setuptools'
overcloud
```

This update replaces the python-setuptools provided by the overcloud-full image with the version provided by the OpenShift repos. Any subsequent yum updates are successful.

BZ#1646707

In some OVS versions, `updelay` and `downdelay` bond settings are ignored, and the default settings are always used.

BZ#1647005

Nova-compute ironic driver tries to update BM node while the node is being cleaned up. The cleaning takes approximately five minutes but nova-compute attempts to update the node for approximately two minutes. After timeout, nova-compute stops and puts nova instance into ERROR state.

As a workaround, set the following configuration option for nova-compute service:

```
[ironic]
api_max_retries = 180
```

As a result, nova-compute continues to attempt to update BM node longer and eventually succeeds.

BZ#1652444

The `neutron_driver` parameter has the value `null` in the `containers-prepare-parameter.yaml` file. This might cause minor updates to the overcloud in OpenDaylight deployments.

Workaround: Before you update the overcloud, set the value of the `neutron_driver` parameter to `odl`.

BZ#1653348

Scaling out with an additional OpenShift master node of a director deployed OpenShift environment fails with a message similar to: "The field 'vars' has an invalid value, which includes an undefined variable. The error was: 'openshift_master_etcd_urls' is undefined..."

BZ#1653466

Scaling out with an additional Infra node on a director deployed OpenShift environment with CNS enabled fails with a message similar to the following: "fatal: [openshift-master-2]: FAILED! => {"changed": false, "msg": "Error mounting /tmp/openshift-glusterfs-registry-c8qImT: Mount failed.""

BZ#1659183

Director and openshift-ansible have different expectations regarding image tags. For example, when importing the remote container images locally, director converts the generic tag into one that uniquely identifies the image based on the `version` and `release` labels from the image metadata. Openshift-ansible, however, relies on a unique `openshift_image_tag` variable for all the openshift images tags making it impossible to specify tags of images individually. Deployment of OCP via director fails when the floating v3.11 tag in the remote container image registry points to images with non-consistent `release` or `version` labels in their metadata.

From the undercloud, import the odd images prior to deploying OpenShift and set the tag to be consistent across all openshift images:

```
skopeo --tls-verify=false copy
docker://registry.access.redhat.com/openshift3/prometheus:v3.11.51-1
docker://192.168.24.1:8787/openshift3/prometheus:v3.11.51-2
```

Deployment of OpenShift from director completes without missing image.

BZ#1660066

Director does not support triggering Red Hat Enterprise Linux OS and OpenShift Container Platform updates on director deployed OpenShift environments. Director deployed OpenShift environments cannot be minor updated.

BZ#1660475

After config-download has generated the playbooks for the Overcloud, if you execute ansible-playbook with --check parameter, it does not work. Expect an error about undefined stdout for ftype. This will be fixed in the next version.

BZ#1664165

Due to a known Ansible issue (<https://github.com/ansible/ansible/issues/24449>) changes to

`/etc/ssh/ssh_known_hosts` are not propagated to existing `nova_compute` and `nova_libvirt` containers within an environment after adding compute hosts.

As a result, live migration, cold migration and instance resize operations using the newly introduced compute hosts will fail as a result as the original hosts will not have the required public SSH keys.

Workaround:

Restarting all `nova_compute` and `nova_libvirt` containers will ensure the updates to `/etc/ssh/ssh_known_hosts` are correctly written.

Detailed steps can be found in the following KCS article:

[OSP14] After successful scale out, live/cold migration and resize fails to the new added compute with 'Host key verification failed'

<https://access.redhat.com/solutions/3792021>

BZ#1664698

A recent change made memory allocation for instances with NUMA topologies pagesize aware. With this change, memory for instances with NUMA topologies can no longer be oversubscribed.

Memory oversubscription is currently disabled for all instances with a NUMA topology, whereas previously only instances with hugepages were not allowed to use oversubscription. This affects instances with an explicit NUMA topology and those with an implicit topology. An instance can have an implicit NUMA topology due to the use of hugepages or CPU pinning.

If possible, avoid the use of explicit NUMA topologies. If CPU pinning is required, resulting in an implicit NUMA topology, there is no workaround.

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#1668219

OpenDaylight was first made available in OSP 13 and is being deprecated in OSP 14.

Our combined OpenDaylight in OpenStack solution will no longer accept new feature enhancements and we would like to inform those who were looking for an OpenDaylight integrated solution from Red Hat to seek alternatives.

OpenDaylight will continue to be supported and receive bug fixes for the duration of the OSP 14 deprecation cycle, with support planned to be completely dropped by the end of the OSP 13 lifecycle (June 27, 2021).

3.2. RED HAT OPENSTACK PLATFORM 14 MAINTENANCE RELEASE 13 MARCH 2019

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

3.2.1. Enhancements

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

BZ#1645489

This enhancement adds the boolean parameter ``NovaLibvirtVolumeUseMultipath``, which provides a value for the multipath configuration parameter ``libvirt/volume_use_multipath`` in the ``nova.conf`` file for Compute nodes. You can set this parameter for each Compute role. Default value is ``False``.

BZ#1658484

This enhancement sets the number of RPC workers to ``1`` by default in OVN tripleo deployments. The goal of this setting is to reduce the number of workers to save memory resources and the number of connections to OVSD, in cases where the Neutron DHCP agent is not deployed alongside OVN services.

BZ#1673172

This enhancement adds the networking-ansible heat parameter ``IronicDefaultNetworkInterface``, which determines the value of the ``default_network_interface`` parameter in the ``ironic.conf`` configuration file. This value is set to the ``neutron`` interface by default, which enables virtual networking through Neutron on bare metal nodes.

Note: The switches attached to the bare metal nodes must be programmable by the networking service if the ``default_network_interface`` is set to ``neutron``.

3.2.2. Known Issues

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

BZ#1691449

There is currently a known issue where director can hang while deploying OCP. This occurs because the fix described in https://bugzilla.redhat.com/show_bug.cgi?id=1671861 is not a part of the ``overcloud-full`` image for the Red Hat OpenStack Platform 14 z1 release. As a workaround, prior to deploying the overcloud, follow the steps below to update the docker package in the ``overcloud-full`` image. For more information on this procedure, see <https://access.redhat.com/articles/1556833>. After completing these steps, you can expect the director to successfully deploy OCP:

```
$ sudo yum install -y libguestfs-tools
```

```
$ virt-customize --selinux-relabel -a overcloud-full.qcow2 --install  
docker  
$ source stackrc  
$ openstack overcloud image upload --update-existing
```

3.3. RED HAT OPENSTACK PLATFORM 14 MAINTENANCE RELEASE 30 APRIL 2019

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. Enhancements

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

[BZ#1658192](#)

```
This feature adds the capability to configure the Cinder Dell EMC  
StorageCenter driver to use a multipath for volume-to-image and image-to-  
volume transfers. The feature includes a new parameter  
`CinderDellScMultipathXfer` with a default value of `True`. Enabling  
multipath transfers can reduce the total time of data transfers between  
volumes and images.
```

3.3.2. Known Issues

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

[BZ#1653348](#)

```
Scaling out with an additional OpenShift master node of a director  
deployed OpenShift environment fails with a message similar to: "The field  
'vars' has an invalid value, which includes an undefined variable. The  
error was: 'openshift_master_etcd_urls' is undefined..."
```

[BZ#1659183](#)

```
Director and openshift-ansible have different expectations regarding image  
tags. For example, when importing the remote container images locally,  
director converts the generic tag into one that uniquely identifies the  
image based on the `version` and `release` labels from the image metadata.  
Openshift-ansible, however, relies on a unique `openshift_image_tag`  
variable for all the openshift images tags making it impossible to specify  
tags of images individually. Deployment of OCP via director fails when the  
floating v3.11 tag in the remote container image registry points to images  
with non-consistent `release` or `version` labels in their metadata.
```

From the undercloud, import the odd images prior to deploying OpenShift and set the tag to be consistent across all openshift images:

```
skopeo --tls-verify=false copy  
docker://registry.access.redhat.com/openshift3/prometheus:v3.11.51-1
```

```
docker://192.168.24.1:8787/openshift3/prometheus:v3.11.51-2
```

Deployment of OpenShift from director completes without missing image.

BZ#1664698

A recent change made memory allocation for instances with NUMA topologies pagesize aware. With this change, memory for instances with NUMA topologies can no longer be oversubscribed.

Memory oversubscription is currently disabled for all instances with a NUMA topology, whereas previously only instances with hugepages were not allowed to use oversubscription. This affects instances with an explicit NUMA topology and those with an implicit topology. An instance can have an implicit NUMA topology due to the use of hugepages or CPU pinning.

If possible, avoid the use of explicit NUMA topologies. If CPU pinning is required, resulting in an implicit NUMA topology, there is no workaround.

BZ#1691449

There is currently a known issue where director can hang while deploying OCP. This occurs because the fix described in https://bugzilla.redhat.com/show_bug.cgi?id=1671861 is not a part of the `overcloud-full` image for the Red Hat OpenStack Platform 14 z1 release. As a workaround, prior to deploying the overcloud, follow the steps below to update the docker package in the `overcloud-full` image. For more information on this procedure, see <https://access.redhat.com/articles/1556833>. After completing these steps, you can expect the director to successfully deploy OCP:

```
$ sudo yum install -y libguestfs-tools
$ virt-customize --selinux-relabel -a overcloud-full.qcow2 --install
docker
$ source stackrc
$ openstack overcloud image upload --update-existing
```

3.3.3. Deprecated Functionality

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

BZ#1687884

As of this release the director graphical user interface is deprecated. Bug fixes and support will be provided through the end of the OSP 13 lifecycle but no new feature enhancements will be made.

CHAPTER 4. TECHNICAL NOTES

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

4.1. RHEA-2019:0045 — OPENSTACK DIRECTOR BUG FIX ADVISORY

The bugs contained in this section are addressed by advisory RHEA-2019:0045. Further information about this advisory is available at link: <https://access.redhat.com/errata/RHEA-2019:0045>

ansible-role-redhat-subscription

BZ#1641180

Previously, the Satellite URL was not correctly set in the role. This prevented the system from getting the Satellite server version, and registration failed. This fix adds the capability to get the `rhsm_satellite_url` value from the `rhsm_baseurl` parameter by default, passes the URL to the registration task to allow force registration, and adds the option to ignore certificate errors. You can override the default value or configure the options as needed.

distribution

BZ#1640095

OpenStack Rally, previously included as a technical preview, is removed from this release.

openstack-aodh

BZ#1467317

With this update, the aodh service now validates event type input queries. Prior to this update, input queries were not validated. An invalid input query could result in the failure to issue an alarm.

openstack-ceilometer

BZ#1521176

Instance resources have three new attributes: `launched_at`, `deleted_at`, `created_at` to track the exact time that Nova creates/launches/deletes instances.

BZ#1596033

The OpenStack Metrics service (Ceilometer) created metrics that were not measured by the Monitoring service (Gnocchi). This fix removes the unnecessary metrics. Now, Ceilometer creates only metrics that will be measured by Gnocchi.

openstack-cinder

BZ#1033180

This release adds a Technology Preview of the ability to attach a volume to multiple hosts or servers simultaneously in both cinder and nova with read/write (RW) mode when this feature is supported by the back end driver. This feature addresses the clustered application workloads use case that typically requires active/active or active/standby scenarios.

BZ#1262068

This enhancement optimizes migration of an RBD volume from one Cinder back end to another when the volume resides within the same Ceph cluster. If both volumes are in the same Ceph cluster, data migration is performed by ceph itself, instead of the cinder-volume process. This reduces migration time.

openstack-ironic

BZ#1649894

Some commands from the OpenStack Bare Metal service (Ironic) to BMCs with IPMI hardware failed due to hardware driver errors. This prevented bare metal nodes from booting. This fix adds the `ipmi_disable_boot_timeout` hardware driver option, which prevents Ironic from sending these commands to IPMI hardware.

BZ#1394888

This update fixes UEFI persistent boot mode support for Dell EMC PowerEdge 13th and 14th generation servers. Those servers now successfully boot into the deployed operating system for either persistent boot modes: BIOS and UEFI.

The fix applies to servers managed by the ironic integrated Dell Remote Access Controller (iDRAC) management hardware implementation ('`idrac`') function, located in `ironic.drivers.modules.drac.management`.

The bug is not resolved for PowerEdge 12th generation and earlier servers; however, BIOS boot mode continues to be supported in PowerEdge 12th generation and earlier servers.

Prior to this update, the boot device would persist during subsequent reboots only when the server's boot mode was set to BIOS.

BZ#1562171

This update introduces multi-tenant bare metal networking with the "neutron" network interface.

By configuring the bare metal nodes with the "neutron" network interface, an operator can enable the users to use isolated VLAN networks for provisioning and tenant traffic on bare metal nodes.

BZ#1638003

In prior releases, a race condition existed in the ironic-conductor hash ring code. A hash ring can be None under load, but this causes an internal server error: 'NoneType' object has no attribute 'getitem_'. *This release fixes the race condition, and ironic API operations no longer fail with 'NoneType' object has no attribute 'getitem_'.*

openstack-keystone**BZ#1462048**

With this update, Users can create application credentials to allow their applications to authenticate to keystone.

See

https://docs.openstack.org/keystone/latest/user/application_credentials.html

openstack-manila-ui**BZ#1600664**

The OpenStack Dashboard (Horizon) plug-in for Manila was unable to retrieve project quota information. This prevented the user from creating shares and caused rendering issues in the shared file systems dashboard. After this fix, the retrieval operation works as expected and the users can view shared file systems and create shares in the dashboard.

openstack-neutron**BZ#1608090**

When using the linuxbridge ml2 driver, non-privileged tenants are able to create and attach ports without specifying an IP address, bypassing IP address validation. A potential Denial of Service could occur if an IP address, conflicting with existing guests or routers, is then assigned from outside of the allowed allocation pool.

openstack-nova**BZ#1512941**

This update supports two new tunable options that can be used to reduce packet drop.

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 update 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#1398343

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

For example, this feature allows to mitigate guests from performance degradation that is caused by running 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 on how to specify granular CPU flags, refer to the documentation of [libvirt]/cpu_model_extra_flags in nova.conf for usage details.

BZ#1402584

With this update, the libvirt compute driver now allows users to create instances with trusted SR-IOV virtual functions. When trusted, a VF can perform certain operations, such as modifying the VF's MAC address in the guest.

Interface bonding requires that all slaves use the same MAC address, which in turn requires MAC address modifications on one of the VFs during a failover. Because MAC address altering is a privileged operation, participating VFs must be trusted in order to successfully configure bonding in the guest.

Administrators can now configure trusted mode for VFs. This requires two steps. First, the 'trusted' value of the '[pci] passthrough_whitelist' JSON configuration option in nova.conf must be set to 'true'. For example:

```
[pci]
passthrough_whitelist = {"devname": "eth0", "trusted": "true",
                        "physical_network": "sriovnet1"}
```

Then, when creating the port, 'trusted=true' must be set for the binding profile. For example:

```
$ neutron port-create <net-id> \
  --name sriov_port \
  --vnic-type direct \
  --binding:profile type=dict trusted=true
```

Because trusted mode only applies to SR-IOV VFs, the 'vnic-type' must be one of 'hw_veb' or 'direct'.

BZ#1469073

Feature:

Add the CPUWeigher weigher for nova-scheduler.

Reason:

The CPUWeigher allows operators to configure a stacking or spreading policy for vCPUs.

Result:

Operators can enable the CPUWeigher and configure a stacking (use all vCPUs on one node first) or spreading (attempt to use a roughly equal amount of vCPUs from all hosts) policy.

BZ#1625122

With this update, Nova screens for the NUMA affinity of host huge pages when booting instances with huge pages. Nova rejects NUMA nodes with insufficient huge pages.

Prior to this update, Nova did not screen for NUMA affinity of huge pages. If the host had insufficient NUMA pages, even with sufficient CPUs, the instance boot would fail.

openstack-sahara**BZ#1547708**

OpenStack Sahara now supports Cloudera Distribution Hadoop (CDH) plugin 5.13.

BZ#1547710

This update adds support of s3-compatible object stores for OpenStack Sahara.

BZ#1639759

A new tripleo heat template has been added to support QDR. Enabling the metrics_qdr service deploys QDR service on all overcloud nodes, which will be used for routing metrics data from collectd service running on each overcloud node.

The relevant heat template parameters are shown below.

MetricsQdrPort:

default: '5666'

description: Service name or port number on which the qdrouterd will accept

connections. This argument must be string, even if the numeric form is used.

type: string

MetricsQdrUsername:

default: 'guest'

description: Username which should be used to authenticate to the deployed

```

        qdrouterd.
type: string
MetricsQdrPassword:
default: 'guest'
description: Password which should be used to authenticate to the deployed
        qdrouterd.
type: string
hidden: true
MetricsQdrConnectors:
default: []
description: Connectors configuration (array of hashes).
type: json
MetricsQdrAddresses:
default:
  - prefix: 'collectd/notify'
    distribution: multicast
  - prefix: 'collectd/telemetry'
    distribution: multicast
description: Addresses configuration (array of hashes).
type: json
MetricsQdrUseSSL:
default: false
description: Set to true if it is required to use SSL or TLS on
        the connection for listener.
type: boolean
MetricsQdrUseEncryption:
default: false
description: Set to true if it is required to encrypt connection to the
        peer
        for listener.
type: boolean
MetricsQdrSaslMechanisms:
default: 'ANONYMOUS'
description: List of accepted SASL auth mechanisms for listener in format
        of comma separated list.
type: string
MetricsQdrSslCertDb:
default: ''
description: Path to SSL certificate db for listener.
type: string
MetricsQdrSslCertFile:
default: ''
description: Path to SSL certificate file for listener.
type: string
MetricsQdrSslKeyFile:
default: ''
description: Path to SSL private key file for listener.
type: string
MetricsQdrSslPwFile:
default: ''
description: Path to SSL password file for certificate key for listener.
type: string
MetricsQdrSslPassword:
default: ''
description: SSL password to be supplied for listener.
type: string

```

```
MetricsQdrTrustedCerts:  
default: ''  
description: Path to file containing trusted certificates for listener.  
type: string
```

BZ#1516911

The `OvsDpdkMemoryChannels` parameter cannot be derived through the DPDK derive parameters workflow. The value is set to 4 by default. You can change that value in your custom environments file to match your hardware.

openstack-tripleo-heat-templates**BZ#1594261**

Prior to this update, with shared storage for `/var/lib/nova/instances`, such as `nfs`, restarting the `nova_compute` container on any compute node resulted in an owner/group change of the instances virtual ephemeral disks and `console.log`. As a result, instances lost access to their virtual ephemeral disks and stopped working.

The method to modify the ownership of the instance files in `/var/lib/nova/instances` have been improved to target only the necessary files/directories.

There is now no loss in access to the instance files during restart of `nova compute`.

BZ#1613847

Dedicated monitor node scale-up or monitor replacement no longer causes the stack update command to fail or to take Ceph monitors out of quorum.

BZ#1637988

After deprecating the `instack_undercloud` functionality, upgrading the undercloud with an admin user failed with a permission error. This was due to the admin user missing the `member` role. This fix adds the `member` role back to the admin user from puppet-keystone module and tripleo-heat-templates.

BZ#1652440

Replacement of Controller nodes with ODL previously failed due to ODL configuration files missing during redeployment. This fix unmounts the `/opt/opendaylight/data` directory from the host, which then triggers the regeneration of ODL configuration files during the replacement process.

BZ#1655151

The OpenStack Director previously configured HAProxy load-balancing with `roundrobin` instead of `source balance`, which resulted in sticky sessions failures. After this fix, the Director uses `source balance` for load-

balancing in the HAProxy configuration, and sticky sessions run as expected.

BZ#1655184

OpenStack Director previously always used IP addresses for the `openshift_master_cluster_hostname` and `openshift_master_cluster_public_hostname` parameters, which caused host names from the `OpenShiftGlobalVariables Heat` parameter to be ignored. After this fix, the Director will use the host name if provided, and the IP addresses if no host name is provided.

BZ#1337770

With this update, OSP 14 supports setting specific IP addresses for each node in each role when using routed spine-and-leaf networking.

Prior to this update, setting specific IPs for each node was only supported for deployments that do not use routed spine-and-leaf networking. It is possible to set specific IPs for each network in OSP 13, but in OSP 13 this feature was considered Tech Preview due to lack of sufficient documentation and testing.

OSP 14 allows operators to choose which IP addresses to use for each network on each node in each role when using routed spine-and-leaf, including multiple routed subnets in the same network.

BZ#1578849

With this update, NTP time is synced early in the deployment process to prevent container configuration and deployment failure. If the NTP servers are not accessible and cannot be synced, deployment fails immediately. Prior to this update, failures could occur later with a cryptic error message.

BZ#1580338

With this enhancement, you can configure `collectd` to send metrics data to local QPID dispatch router services, to support service assurance framework implementation.

New `collectd tripleo heat` template parameters let you configure the `amqp1` write plugin.

Set the `CollectdConnectionType` parameter to `'amqp1'`.

As a result, all metrics data are sent to the local QDR unless explicitly overridden by the parameters shown below.

To deploy the local QDR, use the environment file `/usr/share/openstack-tripleo-heat-templates/environments/metrics-collectd-qdr.yaml`.

The relevant heat template parameters are shown below.

```
CollectdAmqpHost:
type: string
```

```
description: Hostname or IP address of the AMQP 1.0 intermediary.
default: nil
CollectdAmqpPort:
type: string
description: >
    Service name or port number on which the AMQP 1.0 intermediary accepts
    connections. This argument must be a string, even if the numeric form
    is used.
default: '5666'
CollectdAmqpUser:
type: string
description: >
    User part of credentials used to authenticate to the AMQP 1.0
intermediary.
default: guest
CollectdAmqpPassword:
type: string
description: >
    Password part of credentials used to authenticate to the AMQP 1.0
intermediary.
default: guest
hidden: true
CollectdAmqpTransportName:
type: string
description: Name of the AMQP 1.0 transport.
default: metrics
CollectdAmqpAddress:
type: string
description: >
    This option specifies the prefix for the send-to value in the message.
default: collectd
CollectdAmqpInstances:
type: json
description: >
    Hash of hashes. Each inner hash represent Instance block in plugin
    configuration file. Key of outer hash represents instance name.
    The 'address' value concatenated with the 'name' given will be used
    as the send-to address for communications over the messaging link.
default: {}
CollectdAmqpRetryDelay:
type: number
description: >
    When the AMQP 1.0 connection is lost, defines the time in seconds to
wait
    before attempting to reconnect.
default: 1
CollectdAmqpInterval:
type: number
description: >
    Interval on which metrics should be sent to AMQP intermediary. If not
set
    the default for all collectd plugins is used.
default: -666
```

BZ#1613158

In a `tls-everywhere` scenario for VNC, the following TLS connections exist:

- `client` -> `haproxy`
- `novncproxy` -> `vnc server (instance)`

However, the connection from `haproxy` to `nova novncproxy` was not encrypted, resulting in an unencrypted local connection from `haproxy` to `nova novnc-proxy` service on the controller. With this release, the connection from `haproxy` to `nova novnc-proxy` service is encrypted.

BZ#1617927

In prior releases, you could set RX/TX queue size via `nova::compute::libvirt::rx_queue_size/nova::compute::libvirt::tx_queue_size`. However, there was no dedicated TripleO heat template parameter. With this release, the RX/TX queue size can be set on a role base like this:

```
parameter_defaults:
  ComputeParameters:
    NovaLibvirtRxQueueSize: 1024
    NovaLibvirtTxQueueSize: 1024
```

The result is `rx_queue_size/tx_queue_size` is set using new parameters.

BZ#1623387

This update makes `gnocchiclient` available on the undercloud after switching to a containerized undercloud, allowing users to query for telemetry data.

BZ#1635864

Blacklisting configuration updates against Ceph nodes no longer results in failed deployments.

BZ#1640021

Deploying the overcloud with a Gnocchi file backend might fail due to access permissions to the `/var/lib/gnocchi/` directory.

Workaround: Before you deploy the overcloud, set the permissions to the directory in the `openstack/tripleo-heat-templates/docker/services/gnocchi-api.yaml` file as follows:

```
- path:
  list_join:
    - "/"
    - - {get_param: GnocchiFileBasePath}
      - "tmp"
  owner: gnocchi:gnocchi
  perm: '0600'
  recurse: true
```

BZ#1640443

The OpenStack Platform director was not configuring authentication data required for the Block Storage service (cinder) to access privileged portions of the nova API. Because of this, operations on volumes that use nova's privileged API (e.g., migrating an in-use volume) would fail. The director now configures cinder with nova's authentication data. As a result, operations on volumes that require privileges work.

BZ#1652444

The `neutron_driver` is null in the `containers-prepare-parameter.yaml` file. As a result, an OSP14 + ODL minor update will fail if following the general OSP update guide.

Workaround:

Before updating the overcloud, change `neutron_driver` to 'odl' in `containers-prepare-parameter.yaml` (see https://bugzilla.redhat.com/show_bug.cgi?id=1660454).

With this workaround, the OSP14 + ODL minor update should pass.

BZ#1241017

This update adds hostname and network name to the output of the 'openstack port list' command.

The additional information makes it easier to associate Neutron port and IP addresses with a particular host.

BZ#1344174

OSP-Director now supports setting specific IP addresses on the Control Plane (provisioning) network, similar to the way these are set for other networks. To set the IP addresses used during provisioning, set the IPs under the "ctlplane" network, one IP address per line. These will be assigned to the nodes in order (Compute-0, Compute-1, etc.)

For example::

```
parameter_defaults:
  ControllerIPs:
  ctlplane:
    - 192.168.24.251
  ComputeIPs:
  ctlplane:
    - 192.168.24.252
    - 192.168.24.253
```

See the file in `/usr/share/openstack-tripleo-heat-templates/environments/ips-from-pool-ctlplane.yaml` for an example of usage.

Prior to this update, OSP-Director allowed the operator to choose the IP

addresses used by each overcloud node on each network, but the provisioning IP addresses were random. With this update, OSP-Director supports selecting specific IP addresses for each role on every network.

BZ#1477606

With this update, TripleO assigns the default volume type 'tripleo' when creating cinder volumes. Prior to this update, the lack of a volume type caused errors during volume retype and migration operations.

You can change a cinder volume type by overriding the `CinderDefaultVolumeType` parameter.

NOTE: If a cinder default volume type was manually configured (i.e. outside of the Tripleo director), set the `CinderDefaultVolumeType` parameter to the manually configured value when updating the overcloud nodes. This ensures the name of the default volume type doesn't change to the 'tripleo' default value.

BZ#1579866

With this change nova-metadata-api is served via httpd wsgi in the `nova_metadata` container.

Note that upstream will deprecate usage of eventlet for all the WSGI-run services, including nova-api and nova-metadata-api. See <https://review.openstack.org/#/c/549510/> for more details.

BZ#1614282

You can now configure automatic restart of instances on a Compute node if the compute node reboots without first migrating the instances. Nova and the libvirt-guests agent can be configured to gracefully shut down the instances and start them when the Compute node reboots.

New parameters:

`NovaResumeGuestsStateOnHostBoot` (True/False)

`NovaResumeGuestsShutdownTimeout` (default 300s)

BZ#1638922

Previously, the loopback device for Cinder iSCSI/LVM backend was not recreated after a system restart, which prevented the cinder-volume service from restarting. This fix adds a systemd service that recreates the loopback device and therefore persists the Cinder iSCSI/LVM backend after a restart.

openvswitch**BZ#1626488**

A group with no buckets causes Open vSwitch to assert, which results in a daemon crash. With this update, the code allows groups with no buckets. Groups with or without buckets do not trigger the assert.

BZ#1654371

Restarting the service causes internal ports moved to another networking namespace to be recreated. When this happens, the ports lose their networking configuration and are recreated in the wrong networking namespace. With this release, the code does not recreate the ports when the service is restarted, which allows the ports to keep their networking configuration.

BZ#1646707

In some OVS versions, `updelay` and `downdelay` bond settings are ignored, and the default settings are always used.

puppet-nova**BZ#1547954**

With this release, 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 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 `nova.conf` for usage details.

puppet-tripleo**BZ#1614810**

With this update, `logrotate`'s `copytruncate` is used by default for containerized services logs rotation. The default period to keep old logs remains unchanged (14 days).

python-amqp**BZ#1607963**

Cause: When resolving hostnames, `python-amqp` requests both A and AAAA records.

Consequence: If an A record is successfully retrieved, a request for AAAA is still sent. If the resolver is not able to resolve the AAAA record request, the entire operation must wait for the timeout on the AAAA request. This causes name resolution to take longer than necessary, and may cause AMQP messaging operations to be slow or timeout entirely.

Fix: If an A record is successfully returned, do not request the AAAA

record. Simply use the successful A record.

Result: Name lookups resolve quickly.

python-networking-ovn

BZ#1627838

Some versions of OpenStack TripleO Heat templates contained incorrect settings for the Neutron `service_plugins` parameter, which prevented Octavia from working with OVN. This release upgrades the OVN version to support Octavia with the following package: `openstack-tripleo-heat-templates-9.0.1-0`

python-oslo-service

BZ#1642934

Previously, threading events with `eventlet` created unnecessary system calls, which reduced performance of the REST API and resulted in timeout failures in Tempest. This fix improves the response time of the REST API calls, and reduces timeout failures in Tempest.

python-paunch

BZ#1595733

This update corrects an issue that prevented the system from properly shutting down and waiting for containers to stop on reboot. That issue could cause the containers to get killed before they stopped properly. This update adds a new service which ensures that the system waits for the containers to fully stop before continuing during the reboot.

python-pecan

BZ#1597622

Previously, API requests to the policies file for checking non-admin user access permissions caused the entire file to reload and reparse. This resulted in slower processing time and degraded performance.

This bug fix adds caching of the policies file so that queries to the file do not reload the entire file. Now, only changes to the file result in reloading and reparsing the file.

python-tempestconf

BZ#1622011

Rebase package(s) to version:
`python-tempestconf-2.0.0`

In this rebase, the python-tempestconf tool went through a major refactoring to ease the process of generating tempest.conf. This rebase also fixes many outstanding bugs and brings support for new services. Moving to the new release provides many benefits to tempest users. For details, see the OpenStack release notes: <https://docs.openstack.org/releasenotes/python-tempestconf/unreleased.html#relnotes-2-0-0>

python-tripleoclient

BZ#1523328

OpenStack director now uses Ansible for software configuration of the overcloud nodes. Ansible provides a more familiar and debuggable operator experience during overcloud deployment. Ansible is used to replace the communication and transport of the software configuration deployment data between heat and the heat agent (os-collect-config) on the overcloud nodes.

Instead of os-collect-config running on each overcloud node and polling for deployment data from heat, the Ansible control node applies the configuration by running an ansible-playbook with an Ansible inventory file and a set of playbooks and tasks. The Ansible control node (the node running ansible-playbook) is the undercloud by default.

BZ#1601613

The default value of --http-boot changed from /httpboot to /var/lib/ironic/httpboot as containerized Ironic services expect.

BZ#1627041

Previously, sending an IPMI bootdev command caused some hardware to unexpectedly change the boot device order. This prevented some nodes from booting from the correct NIC or prevented PXE from booting from any location.

This release adds a noop management interface for the IPMI driver. This interface handles boot commands and prevents bootdev from being used. To prepare for the noop interface, you must pre-configure nodes to attempt PXE boot mode from the correct NIC, and then fallback to the local hard drive.

python-virtualbmc

BZ#1610505

This update fixes a debug message interpolation bug that caused server crashes when responses were rendered with debug mode activated.

BZ#1624411

During package installation, a bug in the RPM spec for the `virtualbmc` package caused special users or groups that run the `virtualbmc` service to not be created. This update fixes the RPM spec to ensure successful user management operations. The `virtualbmc` service can be successfully started upon package installation.

BZ#1642466

VirtualBMC (VBMC) is no longer supported, and should not be used in production environments. For testing purposes, you can install VBMC directly with `pip`.

openstack-tripleo-common**BZ#1659183**

Director and `openshift-ansible` have different expectations regarding image tags. For example, when importing the remote container images locally, director converts the generic tag into one that uniquely identifies the image based on the version and release labels from the image metadata. `Openshift-ansible`, however, relies on a unique `openshift_image_tag` variable for all the `openshift` images tags making it impossible to specify tags of images individually. Deployment of OCP via director fails when the floating `v3.11` tag in the remote container image registry points to images with non-consistent release or version labels in their metadata.

From the `undercloud`, import the odd images prior to deploying OpenShift and set the tag to be consistent across all `openshift` images:

```
skopeo --tls-verify=false copy
docker://registry.access.redhat.com/openshift3/prometheus:v3.11.51-1
docker://192.168.24.1:8787/openshift3/prometheus:v3.11.51-2
```

Deployment of OpenShift from director completes without missing image.

BZ#1685732

Prior to this update, controllers with a large amount of RAM could experience soft lockups. This would occur due to memory pressure as the dentry cache on controller nodes would grow continually. Now, prior to doing a `curl` statement in the container health check, the `NSS_SDB_USE_CACHE` environment variable is set to `'no'`, which prevents cache growth.

BZ#1660066

Director does not support triggering Red Hat Enterprise Linux OS and OpenShift Container Platform updates on director deployed OpenShift environments. Director deployed OpenShift environments cannot be minor updated.

openstack-tripleo-heat-templates

BZ#1653348

Scaling out with an additional OpenShift master node of a director deployed OpenShift environment fails with a message similar to: "The field 'vars' has an invalid value, which includes an undefined variable. The error was: 'openshift_master_etcd_urls' is undefined..."

BZ#1647956

This update fixes an issue that prevented users from successfully re-running a failed OSP13-to-OSP14 upgrade of OpenStack Platform director. Some upgrade failures resulted in a state where services were not yet deployed with docker, which prevented a successful re-run of the upgrade.

Now a check is performed to verify that the services are deployed under docker control, enabling a successful re-run.

BZ#1652096

This update adds an 'any_errors_fatal' setting to stop an upgrade after an upgradetask failure on any overcloud node.

Prior to this update, after an upgrade failure on one overcloud node, the upgrade would continue on other overcloud nodes.

Now, if an upgrade task fails on any overcloud node, the upgrade is stopped and does not progress onto next tasks on other overcloud nodes.

BZ#1679774

Prior to this update, there was no way to completely disable the Panko service using the Tripleo heat templates.

This is resolved with the newly added parameter, CeilometerEnablePanko. To disable the Panko service, set this parameter to False.

BZ#1654413

Glance image conversion is not enabled by default on a new install of Red Hat OpenStack Platform 14. To use this feature, edit the glance-image-import.conf file.

In the image_import_opts section, insert the following line:
image_import_plugins = ['image_conversion']

BZ#1658192

This feature adds the capability to configure the Cinder Dell EMC StorageCenter driver to use a multipath for volume-to-image and image-to-volume transfers. The feature includes a new parameter CinderDellScMultipathXfer with a default value of True. Enabling multipath transfers can reduce the total time of data transfers between volumes and images.

■

python-os-brick**BZ#1697818**

Previously, the glance-api container image was missing a python library used for managing fibre channel connections. As a result, attempts to create an image on storage would fail when using this medium.

This update provides the python library. The glance-api container now successfully accesses fibre channel storage.