Red Hat OpenStack Platform 16.2

Release Notes

Release details for Red Hat OpenStack Platform 16.2

Last Updated: 2023-11-09

Red Hat OpenStack Platform 16.2 Release Notes

Release details for Red Hat OpenStack Platform 16.2

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

Legal Notice

Abstract

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

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see our CTO Chris Wright's message.

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your input on our documentation. Tell us how we can make it better.

Providing documentation feedback in Jira

Use the Create Issue form to provide feedback on the documentation. The Jira issue will be created in the Red Hat OpenStack Platform Jira project, where you can track the progress of your feedback.

1. Ensure that you are logged in to Jira. If you do not have a Jira account, create an account to submit feedback.
2. Click the following link to open a the **Create Issue** page: Create Issue
3. Complete the **Summary** and **Description** fields. In the **Description** field, include the documentation URL, chapter or section number, and a detailed description of the issue. Do not modify any other fields in the form.
4. Click **Create**.

CHAPTER 1. INTRODUCTION

1.1. ABOUT THIS RELEASE

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

Only changes specific to RHOSP are included in this document. The release notes for the OpenStack "Train" release itself are available at the following location: https://releases.openstack.org/train/index.html.

RHOSP uses components from other Red Hat products. For specific information pertaining to the support of these components, see https://access.redhat.com/site/support/policy/updates/openstack/platform/.

To evaluate RHOSP, sign up at http://www.redhat.com/openstack/.

NOTE

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

1.2. REQUIREMENTS

This version of Red Hat OpenStack Platform (RHOSP) runs on the most recent fully supported release of Red Hat Enterprise Linux 8.4.

The dashboard for this release supports the latest stable versions of the following web browsers:

- Mozilla Firefox
- Mozilla Firefox ESR
- Google Chrome

NOTE

Before you deploy RHOSP, familiarize yourself with the recommended deployment methods. For more information, see Installing and Managing Red Hat OpenStack Platform.

1.3. DEPLOYMENT LIMITS

For a list of deployment limits for Red Hat OpenStack Platform (RHOSP), 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 (RHOSP) 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 (RHOSP), 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 (RHOSP), see Certified Guest Operating Systems in Red Hat OpenStack Platform and Red Hat Enterprise Virtualization.

1.7. PRODUCT CERTIFICATION CATALOG

For a list of the Red Hat Official Product Certification Catalog, see Product Certification Catalog.

1.8. 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 (RHOSP) through Bare Metal Provisioning (ironic), see Supported Operating Systems Deployable With Bare Metal Provisioning (ironic).

1.9. HYPERVISOR SUPPORT

This release of the Red Hat OpenStack Platform (RHOSP) is supported only with the **libvirt** driver (using KVM as the hypervisor on Compute nodes).

This release of the RHOSP runs with Bare Metal Provisioning.

Bare Metal Provisioning has been fully supported since the release of RHOSP 7 (Kilo). You can use Bare Metal Provisioning to provision bare-metal machines by using common technologies such as PXE 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 or non-KVM libvirt hypervisors.

1.10. CONTENT DELIVERY NETWORK (CDN) REPOSITORIES

This section describes the repositories required to deploy Red Hat OpenStack Platform (RHOSP) 16.2.

You can install RHOSP 16.2 through the Content Delivery Network (CDN) using **subscription-manager**. For more information, see Preparing the undercloud.

WARNING

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

1.10.1. Undercloud repositories

RHOSP (RHOSP) 16.2 runs on Red Hat Enterprise Linux (RHEL) 8.4. As a result, you must lock the content from these repositories to the respective RHEL version.

NOTE

- If you synchronize repositories by using Red Hat Satellite, you can enable specific versions of the RHEL repositories. However, the repository label remains the same despite the version you choose. For example, if you enable the 8.4 version of the BaseOS repository, the repository name includes the specific version that you enabled, but the repository label is still **rhel-8-for-x86_64-baseos-eus-rpms**.
- The **advanced-virt-for-rhel-8-x86_64-rpms** and **advanced-virt-for-rhel-8-x86_64-eus-rpms** repositories are no longer required. To disable these repositories, see the Red Hat Knowledgebase solution advanced-virt-for-rhel-8-x86_64-rpms are no longer required in OSP 16.2.

WARNING

Any repositories outside the ones specified here are not supported. Unless recommended, do not enable any other products or repositories outside the ones listed in the following tables or else you might encounter package dependency issues. Do not enable Extra Packages for Enterprise Linux (EPEL).

Core repositories

The following table lists core repositories for installing the undercloud.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for x86_64 - BaseOS (RPMs) Extended Update Support (EUS)	**rhel-8-for-x86_64-baseos-eus-rpms**	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 8 for x86_64 - AppStream (RPMs)	**rhel-8-for-x86_64-appstream-eus-rpms**	Contains RHOSP dependencies.
Red Hat Enterprise Linux 8 for x86_64 - High Availability (RPMs) Extended Update Support (EUS)	**rhel-8-for-x86_64-highavailability-eus-rpms**	High availability tools for RHEL. Used for Controller node high availability.
Red Hat Ansible Engine 2.9 for RHEL 8 x86_64 (RPMs)	**ansible-2.9-for-rhel-8-x86_64-rpms**	Ansible Engine for RHEL. Used to provide the latest version of Ansible.

Name	Repository	Description of requirement
RHOSP 16.2 for RHEL 8 (RPMs)	**openstack-16.2-for-rhel-8-x86_64-rpms**	Core RHOSP repository, which contains packages for RHOSP director.
Red Hat Fast Datapath for RHEL 8 (RPMS)	**fast-datapath-for-rhel-8-x86_64-rpms**	Provides Open vSwitch (OVS) packages for OpenStack Platform.

Ceph repositories

The following table lists Ceph Storage related repositories for the undercloud.

Name	Repository	Description of Requirement
Red Hat Ceph Storage Tools 4 for RHEL 8 x86_64 (RPMs)	**rhceph-4-tools-for-rhel-8-x86_64-rpms**	Provides tools for nodes to communicate with the Ceph Storage cluster. The undercloud requires the **ceph-ansible** package from this repository if you plan to use Ceph Storage in your overcloud or if you want to integrate with an existing Ceph Storage cluster.

IBM POWER repositories

The following table contains a list of repositories for RHOSP on POWER PC architecture. Use these repositories in place of equivalents in the Core repositories.

Name	Repository	Description of requirement
Red Hat Enterprise Linux for IBM Power, little endian - BaseOS (RPMs)	**rhel-8-for-ppc64le-baseos-rpms**	Base operating system repository for ppc64le systems.
Red Hat Enterprise Linux 8 for IBM Power, little endian - AppStream (RPMs)	**rhel-8-for-ppc64le-appstream-rpms**	Contains RHOSP dependencies.
Red Hat Enterprise Linux 8 for IBM Power, little endian - High Availability (RPMs)	**rhel-8-for-ppc64le-highavailability-rpms**	High availability tools for RHEL. Used for Controller node high availability.
Red Hat Fast Datapath for RHEL 8 IBM Power, little endian (RPMS)	**fast-datapath-for-rhel-8-ppc64le-rpms**	Provides Open vSwitch (OVS) packages for OpenStack Platform.

Name	Repository	Description of requirement
Red Hat Ansible Engine 2.9 for RHEL 8 IBM Power, little endian (RPMs)	**ansible-2.9-for-rhel-8-ppc64le-rpms**	Ansible Engine for RHEL. Provides the latest version of Ansible.
Red Hat OpenStack Platform 16.2 for RHEL 8 (RPMs)	**openstack-16.2-for-rhel-8-ppc64le-rpms**	Core RHOSP repository for ppc64le systems.

1.10.2. Overcloud repositories

Red Hat OpenStack Platform (RHOSP) 16.2 runs on Red Hat Enterprise Linux (RHEL) 8.4. As a result, you must lock the content from these repositories to the respective RHEL version.

NOTE

- If you synchronize repositories by using Red Hat Satellite, you can enable specific versions of the RHEL repositories. However, the repository label remains the same despite the version you choose. For example, if you enable the 8.4 version of the BaseOS repository, the repository name includes the specific version that you enabled, but the repository label is still **rhel-8-for-x86_64-baseos-eus-rpms**.
- The **advanced-virt-for-rhel-8-x86_64-rpms** and **advanced-virt-for-rhel-8-x86_64-eus-rpms** repositories are no longer required. To disable these repositories, see the Red Hat Knowledgebase solution advanced-virt-for-rhel-8-x86_64-rpms are no longer required in OSP 16.2.

WARNING

Any repositories outside the ones specified here are not supported. Unless recommended, do not enable any other products or repositories outside the ones listed in the following tables or else you might encounter package dependency issues. Do not enable Extra Packages for Enterprise Linux (EPEL).

Controller node repositories

The following table lists core repositories for Controller nodes in the overcloud.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for x86_64 - BaseOS (RPMs) Extended Update Support (EUS)	**rhel-8-for-x86_64-baseos-eus-rpms**	Base operating system repository for x86_64 systems.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for x86_64 - AppStream (RPMs)	**rhel-8-for-x86_64-appstream-eus-rpms**	Contains RHOSP dependencies.
Red Hat Enterprise Linux 8 for x86_64 - High Availability (RPMs) Extended Update Support (EUS)	**rhel-8-for-x86_64-highavailability-eus-rpms**	High availability tools for RHEL.
Red Hat Ansible Engine 2.9 for RHEL 8 x86_64 (RPMs)	**ansible-2.9-for-rhel-8-x86_64-rpms**	Ansible Engine for RHEL. Used to provide the latest version of Ansible.
Red Hat OpenStack Platform 16.2 for RHEL 8 (RPMs)	**openstack-16.2-for-rhel-8-x86_64-rpms**	Core RHOSP repository.
Red Hat Fast Datapath for RHEL 8 (RPMS)	**fast-datapath-for-rhel-8-x86_64-rpms**	Provides Open vSwitch (OVS) packages for OpenStack Platform.
Red Hat Ceph Storage Tools 4 for RHEL 8 x86_64 (RPMs)	**rhceph-4-tools-for-rhel-8-x86_64-rpms**	Tools for Red Hat Ceph Storage 4 for RHEL 8.

Compute and ComputeHCI node repositories

The following table lists core repositories for Compute and ComputeHCI nodes in the overcloud.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for x86_64 - BaseOS (RPMs) Extended Update Support (EUS)	**rhel-8-for-x86_64-baseos-eus-rpms**	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 8 for x86_64 - AppStream (RPMs)	**rhel-8-for-x86_64-appstream-eus-rpms**	Contains RHOSP dependencies.
Red Hat Enterprise Linux 8 for x86_64 - High Availability (RPMs) Extended Update Support (EUS)	**rhel-8-for-x86_64-highavailability-eus-rpms**	High availability tools for RHEL.
Red Hat Ansible Engine 2.9 for RHEL 8 x86_64 (RPMs)	**ansible-2.9-for-rhel-8-x86_64-rpms**	Ansible Engine for RHEL. Used to provide the latest version of Ansible.
Red Hat OpenStack Platform 16.2 for RHEL 8 (RPMs)	**openstack-16.2-for-rhel-8-x86_64-rpms**	Core RHOSP repository.

Name	Repository	Description of requirement
Red Hat Fast Datapath for RHEL 8 (RPMS)	**fast-datapath-for-rhel-8-x86_64-rpms**	Provides Open vSwitch (OVS) packages for OpenStack Platform.
Red Hat Ceph Storage Tools 4 for RHEL 8 x86_64 (RPMs)	**rhceph-4-tools-for-rhel-8-x86_64-rpms**	Tools for Red Hat Ceph Storage 4 for RHEL 8.

Real Time Compute repositories

The following table lists repositories for Real Time Compute (RTC) functionality.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for x86_64 - Real Time (RPMs)	**rhel-8-for-x86_64-rt-rpms**	Repository for Real Time KVM (RT-KVM). Contains packages to enable the real time kernel. Enable this repository for all Compute nodes targeted for RT-KVM. NOTE: You need a separate subscription to a **Red Hat OpenStack Platform for Real Time** SKU to access this repository.
Red Hat Enterprise Linux 8 for x86_64 - Real Time for NFV (RPMs)	**rhel-8-for-x86_64-nfv-rpms**	Repository for Real Time KVM (RT-KVM) for NFV. Contains packages to enable the real time kernel. Enable this repository for all NFV Compute nodes targeted for RT-KVM. NOTE: You need a separate subscription to a **Red Hat OpenStack Platform for Real Time** SKU to access this repository.

Ceph Storage node repositories

The following table lists Ceph Storage related repositories for the overcloud.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for x86_64 - BaseOS (RPMs)	**rhel-8-for-x86_64-baseos-rpms**	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 8 for x86_64 - AppStream (RPMs)	**rhel-8-for-x86_64-appstream-rpms**	Contains RHOSP dependencies.

Name	Repository	Description of requirement
Red Hat Enterprise Linux 8 for x86_64 - High Availability (RPMs) Extended Update Support (EUS)	**rhel-8-for-x86_64-highavailability-eus-rpms**	High availability tools for RHEL. NOTE: If you used the **overcloud-full** image for your Ceph Storage role, you must enable this repository. Ceph Storage roles should use the **overcloud-minimal** image, which does not require this repository.
Red Hat Ansible Engine 2.9 for RHEL 8 x86_64 (RPMs)	**ansible-2.9-for-rhel-8-x86_64-rpms**	Ansible Engine for RHEL. Used to provide the latest version of Ansible.
Red Hat OpenStack Platform 16.2 Director Deployment Tools for RHEL 8 x86_64 (RPMs)	**openstack-16.2-deployment-tools-for-rhel-8-x86_64-rpms**	Packages to help director configure Ceph Storage nodes. This repository is included with standalone Ceph Storage subscriptions. If you use a combined OpenStack Platform and Ceph Storage subscription, use the **openstack-16.2-for-rhel-8-x86_64-rpms** repository.
Red Hat OpenStack Platform 16.2 for RHEL 8 (RPMs)	**openstack-16.2-for-rhel-8-x86_64-rpms**	Packages to help director configure Ceph Storage nodes. This repository is included with combined OpenStack Platform and Ceph Storage subscriptions. If you use a standalone Ceph Storage subscription, use the **openstack-16.2-deployment-tools-for-rhel-8-x86_64-rpms** repository.
Red Hat Ceph Storage Tools 4 for RHEL 8 x86_64 (RPMs)	**rhceph-4-tools-for-rhel-8-x86_64-rpms**	Provides tools for nodes to communicate with the Ceph Storage cluster.
Red Hat Fast Datapath for RHEL 8 (RPMS)	**fast-datapath-for-rhel-8-x86_64-rpms**	Provides Open vSwitch (OVS) packages for OpenStack Platform. If you are using OVS on Ceph Storage nodes, add this repository to the network interface configuration (NIC) templates.

IBM POWER repositories

The following table lists repositories for RHOSP on POWER PC architecture. Use these repositories in place of equivalents in the Core repositories.

Name	Repository	Description of requirement
Red Hat Enterprise Linux for IBM Power, little endian - BaseOS (RPMs)	**rhel-8-for-ppc64le-baseos-rpms**	Base operating system repository for ppc64le systems.
Red Hat Enterprise Linux 8 for IBM Power, little endian - AppStream (RPMs)	**rhel-8-for-ppc64le-appstream-rpms**	Contains RHOSP dependencies.
Red Hat Enterprise Linux 8 for IBM Power, little endian - High Availability (RPMs)	**rhel-8-for-ppc64le-highavailability-rpms**	High availability tools for RHEL. Used for Controller node high availability.
Red Hat Fast Datapath for RHEL 8 IBM Power, little endian (RPMS)	**fast-datapath-for-rhel-8-ppc64le-rpms**	Provides Open vSwitch (OVS) packages for OpenStack Platform.
Red Hat Ansible Engine 2.9 for RHEL 8 IBM Power, little endian (RPMs)	**ansible-2.9-for-rhel-8-ppc64le-rpms**	Ansible Engine for RHEL. Used to provide the latest version of Ansible.
Red Hat OpenStack Platform 16.2 for RHEL 8 (RPMs)	**openstack-16.2-for-rhel-8-ppc64le-rpms**	Core RHOSP repository for ppc64le systems.

1.11. STREAMLINED CONTAINER BUILD

In Red Hat OpenStack Platform (RHOSP) 16.2 the container build has been streamlined. As a result, several containers have been removed:

- registry.redhat.io/rhosp-rhel8/openstack-glance-base
- registry.redhat.io/rhosp-rhel8/openstack-rsyslog-base
- registry.redhat.io/rhosp-rhel8/openstack-panko-base
- registry.redhat.io/rhosp-rhel8/openstack-keystone-base
- registry.redhat.io/rhosp-rhel8/openstack-openvswitch-base
- registry.redhat.io/rhosp-rhel8/openstack-redis-base
- registry.redhat.io/rhosp-rhel8/openstack-placement-base
- registry.redhat.io/rhosp-rhel8/openstack-zaqar-base

1.12. 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 Red Hat OpenStack Platform (RHOSP) deployment. Facilities available through 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 RHOSP users:

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

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

1.13. UNSUPPORTED FEATURES

The following features are not supported in Red Hat OpenStack Platform (RHOSP):

- Custom policies, which includes modification of **policy.json** files either manually or through any ***Policies** heat parameters. Do not modify the default policies unless the documentation contains explicit instructions to do so.
- Containers are not available for the following packages, therefore they are not supported in RHOSP:
  - **nova-serialproxy**
  - **nova-spicehtml5proxy**
- File injection of personality files to inject user data into virtual machine instances. Instead, cloud users can pass data to their instances by using the **--user-data** option to run a script during instance boot, or set instance metadata by using the **--property** option when launching an instance. For more information, see Creating a customized instance.

If you require support for any of these features, please contact the Red Hat Customer Experience and Engagement team to obtain a support exception.

CHAPTER 2. TOP NEW FEATURES

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

2.1. RED HAT OPENSTACK PLATFORM DIRECTOR

This section outlines the top new features for Red Hat OpenStack Platform (RHOSP) director.

Validation framework output formats

RHOSP contains a validation framework to help verify the requirements and functionality of the undercloud and overcloud. The framework includes new output formats for validation logs:

validation_json

The framework saves JSON-formatted validation results as a log file in **/var/log/validations**. This is the default callback for the validation framework.

validation_stdout

The framework displays JSON-formatted validation results on screen.

http_json

The framework sends JSON-formatted validation results to an external logging server.

Use the **ANSIBLE_STDOUT_CALLBACK** environment variable to set the format that you want with your **openstack tripleo validator run** command:

```
$ openstack tripleo validator run --extra-env-vars ANSIBLE_STDOUT_CALLBACK=<callback> --validation check-ram
```

2.2. BACKUP AND RESTORE

This section outlines the top new features and changes for Red Hat OpenStack Platform (RHOSP) backup and restore components.

Sequential backup for control plane nodes

The backup process for control plane nodes now runs sequentially on each node instead of simultaneously on all nodes. Therefore, you can create a backup of the control plane nodes without service disruption to your environment.

2.3. COMPUTE

This section outlines the top new features for the Red Hat OpenStack Platform (RHOSP) Compute service (nova).

Memory encryption for instances

You can configure AMD SEV Compute nodes to provide cloud users the ability to create instances that use memory encryption. For more information, see Configuring AMD SEV Compute nodes to provide memory encryption for instances.

vGPU resize and cold migration

Instances with a vGPU flavor are automatically re-allocated the vGPU resources after resize and cold migration operations.

Image downloads direct from RBD

You can configure the Compute service to download images directly from the RBD image repository without using the Image service API, when:

- the Image service (glance) uses Red Hat Ceph RADOS Block Device (RBD) as the back end and
- the Compute service uses local file-based ephemeral storage, you can configure the Compute service to download images directly from the RBD image repository without using the Image service API.
  This reduces the time it takes to download an image to the Compute node image cache at instance boot time, which improves instance launch time. For more information, see Configuring image downloads directly from Red Hat Ceph RADOS Block Device (RBD) .

2.4. DISTRIBUTED COMPUTE NODES (DCN)

This section outlines the top new features for Distributed Compute Nodes (DCN) in Red Hat OpenStack Platform (RHOSP).

ML2/OVN support

In RHOSP 16.2, the Modular Layer 2 plug-in with the Open Virtual Network mechanism driver (ML2/OVN) is now fully supported for DCN architectures.

Exclude RAW images from DCN edge sites

In RHOSP 16.2, you can use the **NovaImageTypeExcludeList** with a value of **raw** to exclude raw images from advertisement on edge sites that do not have Ceph storage. Excluding raw images from sites without storage limits the use of unnecessary network and local storage resources.

Externally managed Red Hat Ceph Storage at the edge

With the release of RHOSP 16.2, you can now use Red Hat Ceph Storage that is not deployed by RHOSP director at your edge site.

2.5. NETWORKING

This section outlines the top new features for the Red Hat OpenStack Platform (RHOSP) Networking service.

ML2/OVN support for routed provider networks

Starting in RHOSP 16.2 GA, you can use the Modular Layer 2 plug-in with the Open Virtual Network mechanism driver (ML2/OVN) to deploy routed provider networks. Routed provider networks (RPNs) are common in edge distributed compute node (DCN) and spine-leaf routed data center deployments. RPNs enable a single provider network to represent multiple layer 2 networks (broadcast domains) or network segments, permitting the operator to present only one network to users. For more information, see Deploying routed provider networks in the *Networking Guide*.

Availability zones for ML2/OVS and ML2/OVN

Starting in RHOSP 16.2 GA, with the RHOSP Networking service you can group nodes in availability zones (AZs). For nodes that run crucial services, you can schedule these nodes for resources with high availability. AZs are supported only for the Modular Layer 2 plug-in with the Open Virtual Network (ML2/OVN) and Open vSwitch (ML2/OVS) mechanism drivers. For more information, see Using availability zones to make network resources highly available in the *Networking Guide*.

2.6. STORAGE

This section outlines the top new features for the Red Hat OpenStack Platform (RHOSP) storage services.

Automation for DM-Multipathing redundancy configuration

In RHOSP 16.2.3, the DM-Multipathing redundancy configuration for the Block Storage service (cinder) is now automated.

Sparse image upload

With the Image service (glance) API, you can enable sparse image upload to reduce demand on the image storage back end. In sparse images, the Image service does not interpret null byte (empty) sequences as data, therefore only the data itself consumes storage. This feature is particularly useful in distributed compute node (DCN) environments. Sparse image upload also reduces network traffic and improves the image upload speed.

Multiple back ends

By default, a standard Shared File Systems service (manila) deployment environment file has a single back end. With this release, you can configure the Shared File Systems service to use one or more supported back ends.

Image pre-caching

RHOSP director can pre-cache images as part of the glance-api service. With this release, the image pre-cache feature is fully supported.

Configuring an external NFS share for conversion

The Block Storage service (cinder) can now use an external NFS share to perform image format conversion of Image service (glance) images on the overcloud Controller nodes. Using this functionality prevents the space on the node from being completely filled during a conversion operation.
See Configuring an external NFS share for conversion .

Red Hat OpenShift Container Platform support

The Shared File Systems service (manila) with CephFS through NFS fully supports serving shares to Red Hat OpenShift Container Platform through Manila CSI. This solution is not intended for large scale deployments. For important recommendations, see CephFS NFS Manila-CSI Workload Recommendations for Red Hat OpenStack Platform 16.x.

Support for automating multipath deployments

In release 16.2.4, you can specify the location of your multipath configuration file for your overcloud deployment.

2.7. BARE METAL PROVISIONING

This section outlines the top new features for the Red Hat OpenStack Platform (RHOSP) Bare Metal Provisioning service (ironic).

Policy-based routing

With this enhancement, you can use policy-based routing for RHOSP nodes to configure multiple route tables and routing rules with **os-net-config**. Policy-based routing uses route tables where, on a host with multiple links, you can send traffic through a particular interface depending on the source address. You can also define route rules for each interface.

2.8. NETWORK FUNCTIONS VIRTUALIZATION

This section outlines the top new features for Red Hat OpenStack Platform (RHOSP) Network Functions Virtualization (NFV).

Modify kernel args

RHOSP 16.2 includes an update to allow you to modify the kernel args on a deployed node.

AMD support for SRIOV and DPDK

RHOSP 16.2 includes support for Single Root Input/Output Virtualization(SR-IOV) and Data Plane Development Kit(DPDK) workloads on AMD hosts.

2.9. OTHER FEATURES

Red Hat OpenStack Platform director operator

The Red Hat OpenStack Platform (RHOSP) director operator creates a set of custom resource definitions (CRDs) on top of Red Hat OpenShift Container Platform to manage resources normally created by the RHOSP undercloud. CRDs are split into two types for hardware provisioning and software configuration. The operator includes CRDs to create and manage overcloud nets (IPAM), VMSets (for RHOSP Controllers), and BaremetalSets (for RHOSP Computes). The RHOSP director operator became a fully supported feature shortly after the release of the RHOSP 16.2.4 Maintenance Release, on December 13, 2022.
For more information, see RHBA-2022:8952, Release of containers for Red Hat OpenStack Platform 16.2.4 director operator.

2.10. TECHNOLOGY PREVIEWS

This section provides an overview of the top new technology previews in this release of Red Hat OpenStack Platform (RHOSP).

NOTE

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

Transport Layer Security everywhere (TLS-e) now includes memcached

As a technology preview, you can now configure memcached traffic to be encrypted when you configure TLS-e.

Timemaster (Precision Time Protocol and Chrony)

A technology preview is available that supports the use of **timemaster** to configure Precision Time Protocol (PTP) and Chrony in NFV deployments.

Open vSwitch (OVS) Poll Mode Driver (PMD) Auto Load Balance

You can use Open vSwitch (OVS) Poll Mode Driver (PMD) threads to perform the following tasks for user space context switching:

- Continuous polling of input ports for packets.
- Classifying received packets.
- Executing actions on the packets after classification.
  With this technology preview update, you can modify the following parameters to configure OVS PMD automatic load balance:
- OvsPmdAutoLb
- OvsPmdLoadThreshold

- OvsPmdImprovementThreshold
- OvsPmdRebalInterval
  See Configuring OVS PMD Auto Load Balance .

Security group logging

With this technology preview, you can create packet logs for security groups to monitor traffic flows and attempts into and out of an instance. Each log generates a stream of data about events and appends it to a common log file on the Compute host from which the instance was launched.
You can associate any port of an instance with one or more security groups and define one or more rules for each security group. For example, you can create a rule to allow inbound SSH traffic to any instance in a security group named finance. You can create another rule in the same security group to allow instances in that group to send and respond to ICMP (ping) messages.

Then you can create packet logs to record combinations of packet flow events with the related security groups.

2.11. UPGRADES

This section outlines the top new features for Red Hat OpenStack Platform (RHOSP) upgrades.

Customize base packages after Leapp upgrade

In release 16.2.4, after you upgrade your host from Red Hat Enterprise Linux (RHEL) 7.9 to RHEL 8.4, you can specify additional packages to install in your environment by using the BaseTripeloPackages variable. With this feature, you can customize the base packages that your deployment requires on specific roles. For more information, see Customizing the base packages after a Leapp upgrade .

Upgrade the entire overcloud at once

In release 16.2.4, if you are prepared to take your data plane offline, you can now upgrade the whole overcloud at once. With this enhancement, you complete the upgrade much faster, at the cost of some data plane downtime. For more information, see Speeding up an overcloud upgrade .

Update from any source 16.1.z version to the latest minor Red Hat OpenStack Platform version

Starting in RHOSP 16.2.4, you can update your RHOSP environment from any source 16.1.z version. This enhancement reduces cost and saves time during the update process.

CHAPTER 3. RELEASE INFORMATION

These release notes highlight updates in some or all of the following categories. Consider these updates when you deploy this release of Red Hat OpenStack Platform (RHOSP):

- Bug fixes
- Enhancements
- Technology previews
- Release notes
- Known issues
- Deprecated functionality
- Removed functionality

Notes for updates released during the support lifecycle of this RHOSP release appear in the advisory text associated with each update.

3.1. RED HAT OPENSTACK PLATFORM 16.2.6 MAINTENANCE RELEASE - NOVEMBER 8, 2023

Consider the following updates in Red Hat OpenStack Platform (RHOSP) when you deploy this RHOSP release.

3.1.1. Advisory list

This release of Red Hat OpenStack Platform (RHOSP) includes the following advisories:

RHBA-2023:6307

Red Hat OpenStack Platform 16.2.6 (Train) bug fix and enhancement advisory

RHEA-2023:6230

Red Hat OpenStack Platform 16.2.6 deployment images

RHEA-2023:6229

Red Hat OpenStack Platform 16.2.6 deployment RPM

RHBA-2023:6232

Release of containers for Red Hat OpenStack Platform 16.2.6 (Train)

RHSA-2023:6231

Moderate: Red Hat OpenStack Platform 16.2 (openstack-barbican) security update

3.1.2. Bug fixes

These bugs were fixed in this release of Red Hat OpenStack Platform (RHOSP):

BZ#2004214

This update fixes a NetApp Block Storage (cinder) volume driver issue. Previously, a volume extend operation could fail when the extended size was greater than the maximum LUN geometry on the back end due to a malformed request from the driver. The driver now includes the correct

information in the request.

BZ#2022078

Before this update, you could not suspend and resume instances with mediated (mdev) devices, such as vGPUs, because of a known limitation with the Compute service (nova) and the libvirt driver. This limitation was addressed in the libvirt driver. Now you can suspend and resume instances with mdev devices in RHOSP 16.2, which runs on libvirt 6.0.0 in RHEL 8.4.

BZ#2170683

Before this update, the **db archive** command did not handle large-scale databases efficiently because of the way the command processed child tables. The **db archive** command did not make significant progress, even if you specified a high value for **max_rows** and ran the command repeatedly over long periods of time. With this update, the database archiving code handles large-scale databases with a high number of deleted rows more efficiently, and the **db archive** command makes more predictable progress on large-scale databases.

BZ#2175217

Before this update, the Block Storage API supported the creation of a Block Storage multi-attach volume by passing a parameter in the **volume-create** request, even though this method of creating a multi-attach volume was deprecated for removal. This method can lead to data loss when creating a multi-attach volume on a back end that does not support multi-attach volumes.
The **openstack** and **cinder** CLI only supported creating a multi-attach volume by using a multi-attach volume-type.

With this update, the Block Storage API also only supports creating a multi-attach volume by using a multi-attach volume-type. Some Block Storage API requests that used to work will be rejected with a 400 (Bad Request) response code and an error message.

BZ#2177155

Before this update, in RHOSP environments with ML2/OVN, the north/south traffic for instances with FIPs that belonged to VLAN project (tenant) networks, was centralized instead of distributed, even with Distributed Virtual Routing (DVR) activated. In RHOSP 16.2.6, this issue is fixed in the updated OVN version.

BZ#2179284

Before this update, host services, such as Pacemaker, were mounted under **/var/log/host/** in the rsyslog container. However, the configuration path was the same as the host path **/var/log/pacemaker/**, and the rsyslog service could not retrieve Pacemaker log files. With this update, the Pacemaker log path has been updated to **/var/log/host/pacemaker/**.

BZ#2181566

Before this update, a change in **puppet-tripleo** caused **nova::metadata::dhcp_domain** values in a custom hiera to no longer be applied to Compute hosts. A previously set **dhcp_domain** value became unset, which caused Compute hosts to be renamed. With this update, the goth **nova::metadata::dhcp_domain** and **nova::dhcp_domain** hiera values are honored. The **dhcp_domain** value is preserved across updates, and Compute hosts are not renamed.

BZ#2188051

Before this update, the Block Storage (cinder) Ceph backup driver did not form the internal backup name correctly. As a result, backups that were stored in Ceph could not be restored to volumes that were stored on a non-Ceph backend. With this update, the Red Hat Ceph Storage backup driver forms backup names correctly. Ceph can now identify all the constituent parts of a backup and can restore the data to a volume that is stored on a non-Ceph backend.

BZ#2192413

Before this update, the name of a new networking interface in the amphora instance could conflict

with the name of an existing interface. As a result, adding a new member on a new subnet failed. With this update, the Load-balancing service (octavia) now ensures that the names of the networking interfaces are unique.

BZ#2213409

Before this update, the IPMI agent container did not spawn because the CeilometerIpmi service was not added to THT Compute roles. With this update, the CeilometerIpmi service is added to all THT Compute roles, and the IPMI agent container is spawned with the **--privilege** flag to run **ipmitool** commands on the host. The data collection service (ceilometer) now captures power metrics.

BZ#2217966

Before this update, an incorrect backport resulted in an edge case where images that require additional special type handling might fail to have their headers properly processed. As a result, Red Hat Enterprise Linux CoreOS (RHCOS) 8.6 deployments failed on RHOSP 16.2.
With this update, you can now deploy RHCOS 8.6 based images using RHOSP 16.2. The backport was fixed so that the special pointer handling code can successfully fall back and understand the on-disk format supplied by RHCOS 8.6.

BZ#2233095

On UEFI-based systems, Compute nodes were not booting with the latest kernel version because **/boot/grub2/grubenv** is not a symlink to **/boot/efi/EFI/redhat/grubenv**. With this update, the symlink from **/boot/grub2/grubenv** to **/boot/efi/EFI/redhat/grubenv** is recreated during an update. As a result, the version of the kernel in the **saved_entry** field of **/boot/efi/EFI/redhat/grubenv** is updated, and the Compute node boots with the correct kernel version.

3.1.3. Enhancements

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

BZ#2209090

This update improves the way that the L3 scheduler handles OVN logical router port (LRP) assignments that cannot be correctly completed because there are no gateway chassis available in the availability zone (AZ).
Previously, if you attempted to assign an LRP to an AZ that did not have a gateway chassis, the L3 scheduler incorrectly assigned the LRP to a random set of chassis.

Now the L3 scheduler writes a warning to the logs and assigns the port to a chassis with the warning label “neutron-ovn-invalid-chassis”. There are no changes to existing LRP-to-chassis assignments.

BZ#2240825

This enhancement blocks the Compute service (nova) startup if symptoms of host rename have been detected. The renaming of Compute hosts in a running deployment should never happen, as it has catastrophic consequences on resource tracking and the ability to create new instances or migrate existing ones. Until this enhancement, it was technically possible to rename a Compute host. With this update, the Compute service attempts to detect symptoms of its Compute host getting renamed and does not start if a host rename is detected. This prevents resource tracking corruption and allows the operator to undo the rename before any damage occurs to the deployment. For more information, see Troubleshooting Compute host name change detection.

3.1.4. Technology previews

The items listed in this section are provided as Technology Previews in Red Hat OpenStack Platform (RHOSP). 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#1883298

This update introduces a Technology Preview of the security group logging feature in RHOSP 16.2.6. To monitor traffic flows and attempts into and out of an instance, you can configure Networking service packet logging for security groups.
You can associate any instance port with one or more security groups and define one or more rules for each security group. For example, you can create a rule to drop inbound SSH traffic to any instance in the finance security group. You can create another rule to allow instances in that group to send and respond to ICMP (ping) messages. Then you can configure packet logging to record combinations of accepted and dropped packet flows.

In 16.2, you can use security group logging for stateful security groups. Logged events are stored on the Compute nodes that host the instances, in the file **/var/log/containers/stdouts/ovn_controller.log**.

For more information about Technology Preview features, see Scope of Coverage Details.

For more information about security group logging known issues and workarounds, see https://bugzilla.redhat.com/show_bug.cgi?id=2241184 and https://bugzilla.redhat.com/show_bug.cgi?id=2192918.

3.1.5. Known issues

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

BZ#2008076

Using comma-separated role-specific values for **NeutronBridgeMappings** results in incorrect configuration. **Workaround:** Specify the mappings by using an array in a YAML file instead of a comma-separated value. For example:

```
ComputeParameters:
  NeutronBridgeMappings:
    - datacentre:br-ex
    - datacentre2:br-ex2
```

BZ#2097324

In ML2/OVS, any active connection between two instances or an instance and an external device will not be blocked until the connection is terminated, regardless of whether the security group rule that allows this traffic is removed. **Workaround:** Currently, there is no workaround.

BZ#2192918

A security group logging enhancement introduces a potential race condition that can generate HTTP errors when you create or delete multiple security group logs simultaneously or in rapid succession. The error happens when delete and create requests attempt to delete and create logs in the OVN database at the same time.
Workaround: To avoid this issue, send security group log create and delete requests sequentially instead of in parallel. If there are still errors, add timeouts in between requests.

BZ#2241184

In RHOSP 16.2, an issue in the RHEL implementation of meter bands causes rate and burst limit measurements to be higher than expected for stateful security groups. **Workaround:** If you want fewer security group log entries, decrease the rate and burst limit parameter value in your configuration file.

3.2. RED HAT OPENSTACK PLATFORM 16.2.5 MAINTENANCE RELEASE - APRIL 26, 2023

Consider the following updates in Red Hat OpenStack Platform (RHOSP) 16.2.5 when you deploy this RHOSP release.

3.2.1. Advisory list

This release of Red Hat OpenStack Platform (RHOSP) includes the following advisories:

RHBA-2023:1763

Red Hat OpenStack Platform 16.2 bug fix and enhancement advisory

RHBA-2023:1949

Red Hat OpenStack Platform 16.2.5 (Train) deployment images

RHBA-2023:1950

Red Hat OpenStack Platform 16.2.5 (Train) deployment RPM

RHBA-2023:1951

Release of containers for Red Hat OpenStack Platform 16.2.5 (Train)

RHSA-2023:1948

Low: Red Hat OpenStack Platform 16.2 (openstack-nova) security update

3.2.2. Bug fixes

These bugs were fixed in this release of Red Hat OpenStack Platform:

BZ#1876045

Before this update, endpoints of disabled telemetry services were not cleaned up after an upgrade. This omission did not impact the cloud. With this update, upgrades delete obsolete telemetry endpoints.

BZ#2077944

Before this update, provisioning a network namespace with thousands of subnets took a very long time. This delay prevented the metadata haproxy service from being ready for the first VM started on the hypervisor. As a result, the VM was not properly initialized by the cloud-init process. With this update, improved metadata agent logic for provisioning network namespaces creates faster provisioning performance. This resolves the issue.

BZ#2113819

Before this update, customized Heat policy rules were not applied to the **heat-engine** service. This omission caused some of the customized Heat policy rules defined by the **HeatApiPolicies** paramter to be ignored. With this update, director now generates the customized policy file for all Heat services, including **heat-api**, **heat-api-cfn**, and **heat-engine**. All customized Heat policy rules are now applied. This resolves the issue.

BZ#2133030

Before this update, the Alarming service (aodh) used a deprecated Gnocchi API to aggregate metrics, which sometimes caused Gnocchi to display incorrect CPU use values. With this update, Gnocchi displays the correct metrics by performing calculation, transformation, and aggregation of metrics dynamically.

BZ#2142282

Before this update, the **pure_iscsi_cidr** parameter of the Block Storage service (cinder) PureISCSIDriver did not support IPv6 addresses. With this update, the **pure_iscsi_cidr** parameter of

the Block Storage service PureISCSIDriver supports IPv6 addresses and this driver provides a new parameter called **pure_iscsi_cidr_list**, which supports a list of networks.

BZ#2142684

Before this update, a change in the **auth_encryption_key** parameter caused an inability to delete existing Heat stacks. With this update, Heat allows for changes in the **auth_encryption_key** parameter when deleting existing Heat stacks. Heat ignores objects that cannot be decrypted when deleting Heat stacks.

BZ#2151893

Before this update, inadequate TCP buffer sizes resulted in out of memory warnings for TCP in amphora. The smaller TCP buffer size had a potential negative impact on TCP flows with large payloads. This update increases the size of the TCP buffers in amphora, improving the reliability of the TCP connections. This resolves the issue.

BZ#2153458

Before this update, the Block Storage service (Cinder) miscalculated the amount of free space available on a storage back end after deducting the amount allocated for a newly created volume. These free space calculation errors would accumulate until the next periodic update of the actual amount of free space available on the storage back end. Therefore the rapid creation of multiple volumes could create errors could create errors by falsely indicating that the storage back end was out of space. With this update, Block Storage correctly calculates the available free space on a back end after creating a volume. Therefore multiple volumes can be rapidly created on a back end without errors due to incorrect free space calculations.

BZ#2155987

Before this update, an issue caused glance_api cron jobs from being triggered. With this update, the issue is resolved.

BZ#2159555

Before this update, deploying Block Storage multipath with a custom configuration containing a "blacklist_exceptions" section failed due to an error in the multipath configuration. With this update, custom multipath configurations are handled correctly and therefore Block Storage multipath can be deployed with a custom configuration containing a "blacklist_exceptions" section.

BZ#2165032

Before this update, a race condition occurred in Octavia that might have caused load balancers that use the OVN provider to become stuck in PENDING DELETE under certain conditions. This caused the load balancer to be immutable and unable to update. With this update, the race condition is fixed to resolve the issue.

BZ#2165494

Before this update, the multipath daemon running in a container did not detect changes in the underlying multipath devices on the host. This resulted in failure of Block Storage operations such as resizing an online volume. With this update, the container running the multipath daemon is kept synchronized with multipath devices on the host so that Block Storage operations on multipath volumes function correctly.

BZ#2172897

Before this update, the Block Storage service PowerMax driver failed to properly handle volume snapshots based on the PowerMax legacy snapshot identification method of a generation number. This caused the failure of attempts to manage a snapshot identified by a generation number. With this update, the PowerMax driver properly handles and manages snapshots identified by a generation number.

3.2.3. Enhancements

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

BZ#2148393

With this update, operators can configure the number of metadata agent workers by using the **NeutronMetadataWorkers** parameter in the tripleo heat template. Each OVN metadata agent worker creates a connection to the OVN southbound database. For optimal scaling, avoid overloading the database by setting the worker count to 1.

BZ#2154361

With this update, operators have the ability to configure the number of metadata agent workers using the **NeutronMetadataWorkers** parameter in the THT. Each OVN metadata agent worker creates a connection to the OVN southbound database. For optimal scaling, it is good practice to set the worker count to 1 to avoid overloading the database.

3.2.4. Known issues

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

BZ#2177155

In OpenStack environments with ML2/OVN, the north/south traffic for VMs with FIPs, belonging to VLAN project (tenant) networks, is centralized instead of distributed, even with DVR enabled.
This issue will be fixed in a future release of Fast Datapath (BZ 2007120). Customers experiencing this issue who need an earlier solution should request a hotfix.

BZ#2224236

In this release of Red Hat OpenStack Platform (RHOSP), there is a known issue where SR-IOV interfaces that use Intel X710 and E810 series controller virtual functions (VFs) with the iavf driver can experience network connectivity issues that involve link status flapping. The affected guest kernel versions are:

- RHEL 8.7.0 → 8.7.3 (No fixes planned. End of life.)
- RHEL 8.8.0 → 8.8.2 (Fix planned in version 8.8.3.)
- RHEL 9.2.0 → 9.2.2 (Fix planned in version 9.2.3.)
- Upstream Linux 4.9.0 → 6.4.* (Fix planned in version 6.5.)
  Workaround: There is none, other than to use a non-affected guest kernel.

BZ#2232573

In RHOSP 16.2.5, there is a known issue where SQLAlchemy queries don't retrieve all of the RHOSP Networking service (neutron) RBAC entries for network resources.
For example, if a network has two RBAC registers, one with action **access_as_shared** and other with **access_as_external**, and both are accessible to a specific non-admin user, the **openstack port list --share** query does not show the related network. However, this network is still available to the user because of the **access_as_external** RBAC entry.

Workaround: Currently, there is no workaround.

3.2.5. 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#2195931

If your Red Hat OpenStack Platform (RHOSP) deployment uses Cisco ACI or other third-party ML2 mechanism drivers that rely on **neutron-dhcp-agent**, you must add **DhcpAgentNotification: true** to your custom heat templates before updating to RHOSP 16.2.5.
In RHOSP 16.2.5, DHCP agent notification is now disabled by default. This configuration change facilitates easier deployment of RHOSP with the ML2/OVN mechanism driver, because RHOSP ML2/OVN deployments do not use the DHCP agent. The heat templates used in ML2/OVS deployments were also modified to provide uninterrupted use of DHCP services in ML2/OVS deployments.

If your deployment uses the OVN or OVS ML2 mechanism driver, your updated deployment will not be affected.

Some third-party mechanism drivers use DHCP and rely on the DHCP notifications. If your environment uses one of these third-party mechanism drivers and you update to RHOSP 16.2.5 without enabling DHCP notification (**DhcpAgentNotification: true**), your updated RHOSP deployment is subject to the following issues:

- Network namespaces and **dnsmasq** processes are not created for new networks and ports.
- DHCP and DNS are not served to VM instances on the new networks.
- New instances fail to launch on existing networks with an error message that the instance cannot be scheduled.

3.2.6. Deprecated functionality

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

BZ#2187380

The technology preview support added in RHOSP 16.1 for configuring NVDIMM Compute nodes to provide persistent memory for instances has been deprecated in RHOSP 16.2.5, and will be removed in RHOSP 17.0. Red Hat is removing support for persistent memory from RHOSP 17.0 and future releases in response to the announcement by the Intel Corporation on July 28, 2022 that they are discontinuing investment in their Intel® Optane™ business:

- Intel® Optane™ Business Update: What Does This Mean for Warranty and Support
- Intel® Product Change Notification #119311-00

Cloud operators must ensure that no instances use the vPMEM feature before upgrading to 17.1.

3.3. RED HAT OPENSTACK PLATFORM 16.2.4 MAINTENANCE RELEASE - DECEMBER 7, 2022

Consider the following updates in Red Hat OpenStack Platform (RHOSP) 16.2.4 when you deploy this RHOSP release.

3.3.1. Advisory list

This release of Red Hat OpenStack Platform (RHOSP) includes the following advisories:

RHBA-2022:8794

Release of components for Red Hat OpenStack Platform 16.2.4

RHEA-2022:8842

Red Hat OpenStack Platform 16.2.4 director images

RHEA-2022:8843

Red Hat OpenStack Platform 16.2.4 director image RPMs

RHBA-2022:8844

Updated Red Hat OpenStack Platform 16.2.4 container images

RHSA-2022:8845

Moderate: Red Hat OpenStack Platform 16.2.4 (python-paramiko) security update

RHSA-2022:8846

Moderate: Red Hat OpenStack Platform 16.2.4 (puppet) security update

RHSA-2022:8847

Moderate: Red Hat OpenStack Platform 16.2.4 (protobuf) security update

RHSA-2022:8848

Moderate: Red Hat OpenStack Platform 16.2.4 (python-XStatic-Bootstrap-SCSS) security update

RHSA-2022:8849

Moderate: Red Hat OpenStack Platform 16.2.4 (python-XStatic-Angular) security update

RHSA-2022:8850

Moderate: Red Hat OpenStack Platform 16.2.4 (python-ujson) security update

RHSA-2022:8851

Low: Red Hat OpenStack Platform 16.2 (rabbitmq-server) security update

RHSA-2022:8852

Moderate: Red Hat OpenStack Platform 16.2.4 (numpy) security update

RHSA-2022:8853

Moderate: Red Hat OpenStack Platform 16.2.4 (python-django20) security update

RHSA-2022:8854

Moderate: Red Hat OpenStack Platform 16.2.4 (python-scciclient) security update

RHSA-2022:8855

Moderate: Red Hat OpenStack Platform 16.2.4 (openstack-neutron) security update

RHSA-2022:8856

Low: Red Hat OpenStack Platform 16.2.4 (python-django-horizon) security update

RHSA-2022:8857

Moderate: Red Hat OpenStack Platform 16.2.4 (erlang) security update

3.3.2. Bug fixes

These bugs were fixed in this release of Red Hat OpenStack Platform:

BZ#1942717

This release supports port filtering for the Dell EMC XtremeIO driver for the Block Storage service (cinder).

BZ#2057002

Before this update, a race condition occurred when the Compute service (nova) requested that the Block Storage service (cinder) detach a volume and then an external request was made to delete this same volume. This could result in the volume being deleted first before it was detached, which prevented the Compute service from removing this non-existent volume. With this update, if the Compute service requests that the Block Storage service detach a volume and then an external request is made to delete this same volume, this volume is always detached first and then it is deleted.

BZ#2092088

This update fixes a bug that prevented the ceilometer-agent-compute service from collecting libvirt-related metrics.
Previously, the libvirt service started after the ceilometer-agent-compute service, which resulted in "Permission denied" failures and loss of metrics data. Now the libvirt service starts before the ceilometer-agent-compute service and the service can properly collect metrics.

BZ#2094377

Previously, Red Hat Ceph Storage nodes were incorrectly configured to consume OpenStack high availability, advanced-virt, and fast-datapath repos during Leapp upgrades. The previous bug fix for this issue introduced an override that caused role-based parameters to work incorrectly.
With this update, the role-based parameter implementation is fixed and the correct repositories are enabled for Red Hat Ceph Storage nodes. This update fixes the issue in Red Hat OpenStack Platform environments 16.2 and later that use the Red Hat Ceph Storage role.

BZ#2103970

This update fixes a bug that prevented the ceilometer-agent-compute service from polling for CPU metrics on Compute nodes.

BZ#2106647

Before this update, in overcloud deployments that enabled the Block Storage (cinder) backup service, a stack update affecting the Block Storage configuration did not restart the Block Storage service. This caused the Block Storage service to use the old configuration. With this update, the stack update procedure ensures that both the Block Storage backup service and the Block Storage service restart when the Block Storage configuration changes. This ensures that the Block Storage service always uses the latest configuration.

BZ#2109350

This RHOSP 16.2.4 update makes it possible for you to correct a libvirt version incompatibility before updating to RHOSP 16.2.4. If your deployment has the incompatibility issue and you do not perform the steps published in the KCS article workaround before updating to RHOSP 16.2.4, the update might leave instances in an unmanageable state.
Before updating to 16.2.4, see the KCS article Workaround for a libvirt version-compat issue (bug 2109350) when updating RHOSP 16.2.0.

Perform the steps in the article to determine whether your update path is affected by the libvirt incompatibility issue. If it is affected, perform the remaining steps in the KCS article to resolve the issue.

BZ#2111871

This update fixes a bug that causes connectivity loss after certain updates to RHOSP 16.2.2 and 16.2.3. If you are planning to update to a RHOSP 16.2 release, update to RHOSP 16.2.4 to avoid connectivity loss.
The bug is triggered by a database schema change in OVN 21.12, which is introduced in RHOSP 16.2.2. and 16.2.3. OVN 21.12 contains a new column that is not present in earlier versions. OVN database schema changes should not cause a problem in OpenStack, but this particular change is affected by a bug.

In particular, instance connectivity is lost for a variable amount of time (from 20 seconds to 3 minutes) when you run the following command:

```
$ openstack overcloud external-update run --stack overcloud --tags ovn
```

To avoid the bug, do not update to RHOSP 16.2.2. or 16.2.3. Update to RHOSP 16.2.4 instead.

BZ#2112918

Before this update, the oslo-config-validation falsely reported errors with the "key_manager" and "barbican" sections of Block Storage (cinder) configuration. With this update, oslo-config-validation no longer falsely reports Block Storage configuration errors.

BZ#2119145

Before this update, when the number of objects in the Object Storage service (swift) container for the overcloud exceeded 10,000, only the first 10,000 objects in the config-download directory were cleaned up during a delete operation, and the remaining objects prevented the container from getting deleted. With this update, there is added handling for cases where there are more than 10,000 objects in the Object Storage service container for the overcloud.

BZ#2123226

Before this update, VM instances (amphorae) for the Red Hat OpenStack Platform (RHOSP) Load-balancing service (octavia) could experience performance issues when a lot of connections filled the network connection tracking (conntrack) table. The cause for this was that conntrack was enabled for all packet types, including TCP, which does not require conntrack. In RHOSP 16.2.4, amphora performance has improved, because conntrack is disabled for TCP packets and is only enabled for UDP and SCTP packets.

BZ#2123318

Before this update, an SELinux issue triggered errors with Red Hat OpenStack Platform (RHOSP) Load-balancing service (octavia) ICMP health monitors that used the amphora provider driver. In RHOSP 16.2.4, this issue has been fixed and ICMP health monitors function properly.

BZ#2126616

This update fixes the following PowerMax Block Storage (cinder) driver issues:

- Before this update, the PowerMax Block Storage driver deleted all non-temporary snapshots during the **do_sync_check** operation. This update adds a check to determine if a snapshot must be deleted. This ensures that the **do_sync_check** operation does not indiscriminately delete non-temporary snapshots.

- Before this update, the PowerMax Block Storage driver used case-sensitive conditions which could return errors when modifying the storage group. With this update, these conditions are case-insensitive and storage groups can be modified successfully.

BZ#2126786

This update fixes a bug that prevented the Telemetry service (ceilometer) from polling the Object Storage service (swift) for metrics. The Telemetry service now polls the Object Storage service correctly.

BZ#2131386

This update fixes a bug that deleted the existing Block Storage (cinder) backup record when a backup record was imported for an existing **backup_id**.

BZ#2134529

Deploying RHEL 8.6 images in UEFI mode failed when using the ironic-python-agent, because the ironic-python-agent service did not understand the RHEL 8.6 UEFI boot loader hint file. With this update, you can now deploy RHEL 8.6 in UEFI mode.

BZ#2137484

RHSA-2022:6969 introduced the process to clean up files in the /var/lib/mistral directory in the undercloud but the process consistently failed when the Load-balancing service (octavia) or Red Hat Ceph Storage was enabled because these services created additional directories, which the cleanup process could not properly remove. Some deployment actions, such as scale out, consistently failed if the Load-balancing service or Ceph Storage was enabled. With this update, Mistral no longer executes the cleanup. Users must manually delete files if they want to enforce the reduced permission of the files in the /var/lib/mistral directory. Deployment actions no longer fail because of a permission error.

BZ#2138203

Before this update, some deployment or scale-up operations with affected builds failed due to a missing **tuned-profiles-cpu-partitioning** package on nodes that used the **overcloud-minimal** baremetal image, for example, Red Hat Ceph Storage nodes. With this update, the **tuned-profiles-cpu-partitioning** package is included in the **overcloud-minimal** image. This update restores the deployment and scale-up functionality for nodes that use the **overcloud-minimal** image.

BZ#2138395

This update fixes a bug that caused ceilometer-agent-ipmi to write log data inside the container namespace instead of on the host as expected.
The improper placement of the content inside the container increased the container size, prevented proper log rotation, and resulted in loss of the log data when the container was deleted or rebuilt.

Now ceilometer-agent-ipmi writes the logs to the host in **/var/log/containers/ceilometer/** as expected.

3.3.3. Enhancements

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

BZ#1933751

This enhancement adds a method for pulling down third-party containers by introducing a Jinja template processing approach and adding a **template basedir** parameter, which is required by the Jinja loader, to the BaseImageManager. With this update, pulling down the Ceph-related containers is now optional. You can avoid pulling down the Ceph-related containers by setting the **ceph_images** Boolean value to **False**.

BZ#1990357

This enhancement adds new configuration options for the Networking service (neutron) logging service plug-in. With this update, there is added support for network security group logging, and you can configure Networking service logging by using the following new parameters in RHOSP Orchestration service (heat) templates:

- **Options for Layer 3 (L3) agents:**
  - **NeutronL3AgentLoggingRateLimit**
  - **NeutronL3AgentLoggingBurstLimit**
  - **NeutronL3AgentLoggingLocalOutputLogBase**
- **Options for for Open vSwitch (OVS) agents:**

- **NeutronOVSAgentLoggingRateLimit**
- **NeutronOVSAgentLoggingBurstLimit**
- **NeutronOVSAgentLoggingLocalOutputLogBase**

- Options for ML2/OVN back ends:
  - **NeutronOVNLoggingRateLimit**
  - **NeutronOVNLoggingBurstLimit**
  - **NeutronOVNLoggingLocalOutputLogBase**

BZ#2027851

If you are prepared to take your data plane offline, you can now upgrade the whole overcloud at once. With this enhancement, you complete the upgrade much faster, at the cost of some data plane downtime. For more information, see Speeding up an overcloud upgrade .

BZ#2037996

You can now specify a different Controller group name than the default, **Controller,** by using the `controller_group_name ` variable.

BZ#2075039

With this update, you can now migrate an ML2/OVS deployment with the iptables_hybrid firewall driver to ML2/OVN.

BZ#2081630

Starting in Red Hat OpenStack Platform (RHOSP) 16.2.4, you can update your RHOSP environment from any source 16.1.z version. This enhancement reduces cost and saves time during the update process.

BZ#2102017

With this update, you can use Distributed Virtual Routing (DVR) to route traffic to VLAN project networks in an ML2/OVN deployment.
Previously, if you had VLAN tenant networks and DVR in an ML2/OVN deployment, the Networking service routed traffic centrally. Now, with DVR enabled, traffic routed to the VLAN networks goes directly to a node hosting ports.

BZ#2123646

After you upgrade your host from Red Hat Enterprise Linux (RHEL) 7.9 to RHEL 8.4, you can specify additional packages to install in your environment by using the BaseTripeloPackages variable. With this feature, you can customize the base packages that your deployment requires on specific roles. For more information, see Customizing the base packages after a Leapp upgrade .

3.3.4. 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#1992655

In previous releases, administrators had to add the **ceph** plugin to the **CollectdExtraPlugins** parameter in their custom environment files.

With this release, the **ceph** plugin loads automatically on Ceph Storage nodes. Therefore, before you upgrade from Red Hat OpenStack Platform 13 to 16.2, you must remove the **ceph** plugin from the **CollectdExtraPlugins** parameter in your custom environment files.

BZ#2022940

Director support has been added to configure the Shared File Systems service (manila) with Dell EMC PowerMax storage systems.

3.3.5. Known issues

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

BZ#2097444

There is currently a known issue that the OVS minimum bandwidth value is not cleared from the port when a user removes the network policy. The workaround is to manually delete the Queue register from the OVS database:

```
$ ovs-vsctl destroy Queue <queue_uuid>
```

You can locate the Queue register by looking for the Queue **external_ids:port** reference, that contains the Neutron port ID.

BZ#2134557

You can cause framing errors if you configure an ID value longer than 62 characters for the **metrics_qdr** service. An example error message is **failed: amqp:connection:framing-error connection aborted**. When the **metrics_qdr** service is unstable, no telemetry data flows to Service Telemetry Framework (STF).
Workaround: Do not set the **metrics_qdr** ID value longer than 62 characters. The default value for the router ID is **Router.<fqdn>**, where **<fqdn>** is the fully-qualified domain name of the node.

3.3.6. Removed functionality

BZ#2101944

With this update, the collectd processes plugin has been removed from the default list of plugins. Loading the plugin can cause flooding issues and does not provide value when running in a containerized environment as it will only see the collectd and sensubility processes rather than the expected system processes.

3.4. RED HAT OPENSTACK PLATFORM 16.2.3 MAINTENANCE RELEASE - JUNE 22, 2022

Consider the following updates in Red Hat OpenStack Platform (RHOSP) 16.2.3 when you deploy this RHOSP release.

3.4.1. Advisory list

This release of Red Hat OpenStack Platform (RHOSP) includes the following advisories:

RHBA-2022:4793

Release of components for Red Hat OpenStack Platform 16.2.3 (Train)

RHBA-2022:5117

Red Hat OpenStack Platform 16.2.3 containers bug fix advisory

RHEA-2022:5113

Red Hat OpenStack Platform 16.2.1 (Train) director images enhancement advisory

RHSA-2022:5114

Moderate: Red Hat OpenStack Platform 16.2 (openstack-barbican) security update

RHSA-2022:5115

Moderate: Red Hat OpenStack Platform 16.2 (python-django20) security update

RHSA-2022:5116

Moderate: Red Hat OpenStack Platform 16.2 (puppet-firewall) security update

3.4.2. Bug fixes

These bugs were fixed in this release of Red Hat OpenStack Platform:

BZ#1678328

If you create a share without specifying a default share type, the share creation fails. Before this update, the share creation failed without an error message. With this update, an error message is displayed to help you address the reason for the failure.

BZ#2010990

Before this update, manual attempts to trigger introspection of a bare metal node that had previously failed would return the following error message:
"Can not transition from state 'uninitialized' on event 'sync'".

This was due to a defect in the Bare Metal Provisioning inspector (**ironic-inspector**) internal state tracking. The Bare Metal Provisioning inspector was blocking attempts to re-introspect a previously failed bare metal node because the internal state cache indicated that the session was already underway. With this update, new introspection requests explicitly reset the state cache if an entry already exists, and introspection operations that have previously failed can now be retried.

BZ#2015699

Before this update, the **RsyslogElasticsearchSetting** parameter did not support the array format required to set multiple servers. With this update, the **RsyslogElasticsearchSetting** parameter now accepts values with multiple items. You can define multiple Elasticsearch servers by using the following array format: **server: ["elasticsearch1", "elasticsearch2"]**

BZ#2038881

Libvirt was updated and as a result provides more metrics. Due to API changes, collectd was incompatible, resulting in a potential application crash. With this update, collectd was adjusted to provide hugepage usage via the virt plugin. The hugepages metric is exposed via the virt plugin and collectd was adjusted to match the API changes so that the application no longer crashes when pulling the virt metrics.

BZ#2055409

Before this update, during the replacement of a DCN node, the etcd service on the replacement node failed to start and caused the cinder-volume service on that node to fail. This failure was caused by the replacement for a DCN node attempting to start the etcd service as if it were bootstrapping a new etcd cluster, instead of joining the existing etcd cluster.
With this update, a new parameter has been added, **EtcdInitialClusterState**. When **EtcdInitialClusterState** is set to **existing**, the DCN node starts etcd correctly, which causes the cinder-volume service to run successfully.

BZ#2056918

Before this update, attempts to stop the cinder-volume service running in active-active mode resulted in the **failed** state. The cause for these failures was that the service was not allowing sufficient time to properly stop itself.
With this update, the time period allocated to stopping the service has been extended, and the service moves to the **inactive** state when you stop it.

BZ#2064767

Before this update, the RHOSP Bare Metal service (ironic) could lose its connection to the remote Redfish baseboard management controller (BMC) resulting in the bare metal node entering a maintenance state and with its power status changing to **None**. Depending on environmental factors for a site, some or all of the bare metal nodes could be in this unwanted maintenance state for an extended period of time.
Transient network connectivity issues caused by high packet loss to the BMC caused connection caching issues when using Redfish. In cases where a session token needed to be renegotiated, the cached session object was never invalidated and connectivity was lost to the BMC.

With this update, the Bare Metal service now initializes an entirely new cached session with a remote Redfish BMC when connectivity or authentication issues are detected. Additionally, this enables you to use updated credentials if the BMC passwords for the nodes are changed in the future.

BZ#2073096

Before this update, if the **CinderDefaultVolumeType** parameter contained a non-alphanumeric character, the parameter validation failed. As a result, overcloud deployments and updates to existing deployments failed. With this update, **CinderDefaultVolumeType** validation supports non-alphanumeric characters.

3.4.3. Enhancements

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

BZ#1868940

This enhancement adds the type **HostDomain**. **HostDomain** is the same as **HostAddress** with the added support of the underscore character - RFC1033. Systems such as DomainKeys and service records use the underscore. The Compute service can use the **HostDomain** type to define **live_migration_inbound_addr**.

BZ#1892796

In Red Hat OpenStack Platform (RHOSP) 16.2.3, Intel Columbiaville E810 NICs in NFV deployments are now fully supported with the following known issues:

- OVS is only able to use one VF. (BZ#1978622)
- PMD cores running in kernel space for OVN-DPDK. (BZ#2049442)

BZ#1916092

Starting in Red Hat OpenStack Platform (RHOSP) 16.2.3, the Modular Layer 2 mechanism driver with Open Virtual Networking (OVN) supports the VXLAN tunneling protocol. You can now migrate from ML2/OVS to ML2/OVN and continue using VXLAN tunneling. For more information, see the Migrating the Networking Service to the ML2/OVN Mechanism Driver guide.

BZ#1939563

You can now use the **OVNEncapType** option in TripleO Heat templates. With this enhancement, you can set the VXLAN tunnel protocol for the Networking service (neutron), instead of the default,

Geneve. When you specify VXLAN in the **OVNEncapType** option, Open Virtual Network (OVN) uses VXLAN for OpenStack Networking tenant networks.

BZ#1967680

In Red Hat OpenStack Platform (RHOSP) 16.2.3, there is RHOSP Orchestration service (heat) support for Dynamic Device Personalization (DDP) used in Intel Ethernet 800 Series NICs, such as the Columbiaville E810. For more information, see the Network Functions Virtualization Planning and Configuration Guide.

BZ#2042143

In Red Hat OpenStack Platform (RHOSP) 16.2.3, there is support for upgrading firmware and configuring NVIDIA Mellanox BlueField-2 into ConnectX mode by using the mstflint tool, with these two known issues:

- If your RHOSP deployment uses **os-net-config-mappings.yaml** for NIC ordering, then you must use a custom first-boot.yaml file.
- Set tripleo cloud init timeout through templates. (BZ#2097271)

BZ#2047705

In Red Hat OpenStack Platform (RHOSP) 16.2.3, the DM-Multipathing redundancy configuration for the Block Storage service (cinder) is now automated.

BZ#2061697

This enhancement adds a new heat parameter, **FsAioMaxNumber**, that limits the number of parallel AIO requests to 1048576 by default.

3.4.4. Technology previews

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#1846101

In Red Hat OpenStack Platform (RHOSP) 16.2.3, a technology preview is available that supports Open vSwitch (OVS) Traffic Control (TC) Flower classifier hardware offload using connection tracking (conntrack), with the following known issue:

- OVS minimum bandwidth is not cleaned when network policy is removed from port. (BZ#2097444)

3.4.5. 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#2081357

With RHOSP 16.2.3, fencing agents that depend on potentially insecure protocols, such as Telnet, have been removed from the default overcloud images. Overcloud images now include the following fencing agents by default:

- fence-agents-redfish
- fence-agents-ipmilan

- fence-agents-kdump
- fence-agents-rhevm
- fence-agents-compute
- sbd
- fence-agents-sbd

BZ#2092946

Open vSwitch (OVS) does not support offloading OpenFlow rules that have the **skb_priority**, **skb_mark**, or output queue fields set. Those fields are needed to provide quality-of-service (QoS) support for virtio ports.
If you set a minimum bandwidth rule for a virtio port, the Neutron Open vSwitch agent marks the traffic of this port with a Packet Mark Field. As a result, this traffic cannot be offloaded, and it affects the traffic in other ports. If you set a bandwidth limit rule, all traffic is marked with the default 0 queue, which means no traffic can be offloaded.

As a workaround, if your environment includes OVS hardware offload ports, disable the packet marking in the nodes that require hardware offloading. After you disable the packet marking, it will not be possible to set rate limiting rules for virtio ports. However, differentiated services code point (DSCP) marking rules will still be available.

In the configuration file, set the **disable_packet_marking** flag to **true**. After you edit the configuration file, you must restart the **neutron_ovs_agent** container. For example:

```
$ cat `/var/lib/config-data/puppet-
generated/neutron/etc/neutron/plugins/ml2/openvswitch_agent.ini`
  [ovs]
  disable_packet_marking=True
```

3.4.6. Known issues

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

BZ#1986423

Rebooting a node with a virtual function (VF) attached to OVS-DPDK (vfio-pci driver) results in VF uninitialized on that physical function (PF). As a result, virtual machines are unable to use the VFs from that PF. If a second VF is used for another OSP network, it does not function as expected after reboot.
As a workaround, perform the following steps on the Compute node before you reboot the node:

1. Delete the file **/etc/udev/rules.d/70-os-net-config-sriov.rules**.
2. Modify the **Before** criteria of **/etc/systemd/system/sriov_config.service** file to add **network-pre.target**. The modified **Before** should look like:

   ```
   Before=network-pre.target openvswitch.service
   ```

   The workaround fixes the issue and all the VFs initialize correctly.

BZ#2050765

There is currently a known issue in which the OpenStack **tripleo validator run** command produces errors. In the command output the Status_by_Host column contains the value, 'No host matched'. This error occurs with all **openstack validator run** calls because no ansible inventory was available for the undercloud and is not limited to a single validator group. This error also impacts overcloud validations. The root cause of this is a regression in the tripleo client code where tripleo-ansible-inventory is no longer called in real time.
As a result, in updates to Red Hat OpenStack Platform (RHOSP) 16.2, the **tripleo validator run --group pre-introspection** command fails on all tests.

Workaround: create a file called inventory.yaml by manually running **tripleo-ansible-inventory --static-yaml-inventory inventory.yaml**. Then, run the validation using the **-i inventory.yaml** argument:

```
# tripleo-ansible-inventory --static-yaml-inventory inventory.yaml
# openstack tripleo validator run  --group pre-introspection -i inventory.yaml
```

For more information, see the Red Hat Knowledgebase solution openstack tripleo validator run command produces error.

BZ#2097444

There is currently a known issue that the OVS minimum bandwidth value is not cleared from the port when a user removes the network policy. The workaround is to manually delete the Queue register from the OVS database:

```
$ ovs-vsctl destroy Queue <queue_uuid>
```

You can locate the Queue register by looking for the Queue **external_ids:port** reference, that contains the Neutron port ID.

3.5. RED HAT OPENSTACK PLATFORM 16.2.2 MAINTENANCE RELEASE - MARCH 23, 2022

Consider the following updates in Red Hat OpenStack Platform (RHOSP) 16.2.2 when you deploy this RHOSP release.

3.5.1. Advisory list

This release of Red Hat OpenStack Platform (RHOSP) includes the following advisories:

RHEA-2022:1002

Red Hat OpenStack Platform 16.2.2 (Train) deployment images bug fix advisory

RHBA-2022:1001

Release of components for Red Hat OpenStack Platform 16.2.2

RHSA-2022:1000

Moderate: Red Hat OpenStack Platform 16.2 (numpy) security update

RHSA-2022:0999

Moderate: Red Hat OpenStack Platform 16.2 (openstack-nova) security update

RHSA-2022:0998

Moderate: Red Hat OpenStack Platform 16.2 (golang-github-vbatts-tar-split) security update

RHSA-2022:0997

Moderate: Red Hat OpenStack Platform 16.2 (golang-qpid-apache) security update

RHSA-2022:0996

Moderate: Red Hat OpenStack Platform 16.2 (openstack-neutron) security update

RHSA-2022:0995

Moderate: Red Hat OpenStack Platform 16.2 (openstack-tripleo-heat-templates) security update

RHBA-2022:0994

Release of Red Hat OpenStack Platform 16.2 container images

RHSA-2022:0993

Moderate: Red Hat OpenStack Platform 16.2 (python-oslo-utils) security update

RHSA-2022:0992

Important: Red Hat OpenStack Platform 16.2 (python-twisted) security update

3.5.2. Bug fixes

These bugs were fixed in this release of Red Hat OpenStack Platform:

BZ#1956785

Before this update, you could only configure the **net_config_override** parameter in **undercloud.conf** to point to a file in JSON format. With this update, you can now configure the **net_config_override** parameter in **undercloud.conf** to point to a file in either YAML or JSON format.

BZ#1961237

Before this update, the dnsmasq healthcheck failed even when dnsmasq ran correctly. The healthcheck failed because it used the dnsmasq user rather than the root user, and did not have access to the /**proc** files. This resulted in incorrect systemd journal messages and failures when validations were enabled. With this update, the dnsmasq healthcheck is disabled because it is of limited use and it is being phased out in later releases. The dnsmasq container is now marked as healthy as long as it is running.

BZ#1978228

Before this update, the leapp process failed when the following statements were true:

- The system is deployed with TLS-everywhere (TLSe).
- The system uses the deprecated **authconfig** utility to configure authentication on your system.
- The command "leapp answer --section authselect_check.confirm=True" has not been run before triggering the leapp upgrade.
  With this update, the leapp process completes successfully without the need to run the command "leapp answer --section authselect_check.confirm=True" before triggering the leapp upgrade command.

BZ#2003762

Before this update, when creating a snapshot with PowerMaxOS 5978.711, REST experienced a payload response change and caused the device label to modify its format. The underlying data from the solutions enabler changed and no longer contained a colon character (:). This caused an IndexError exception in the PowerMax Driver:

```
IndexError: list index out of range
```

With this update, the problem is resolved in PowerMaxOS 5978.711 and later.

BZ#2006556

This update fixes a bug that omitted details from the output of the **openstack volume backup list** command when the output exceeded 1000 lines.

BZ#2006970

In cases where high CPU use was monitored in a multi-core system, the calculation for CPU use was inaccurate.
With this update, the calculation of CPU use in a multi-core scenario is now accurate. The latest STF dashboards have been adjusted to incorporate this update.

BZ#2010246

This update fixes a bug that caused unintended results when converting Dell EMC PERC H755 RAID controller physical disks to non-RAID mode.
The conversion erroneously created RAID-0 virtual disks and moved them to the **Online** state, consuming a physical disk.

The RAID-0 virtual disks are no longer created during the conversions.

BZ#2022018

This update fixes a bug that caused some virtual machine (VM) bootup operations to fail when multiple VMs were booted simultaneously from an image.
Previously, the Block Storage service (cinder) GPFS SpectrumScale driver did not correctly detect when the storage back end supported copy-on-write (COW) mode. The driver disabled COW features such as the ability to rapidly create volumes from an image. When booting multiple instances simultaneously from an image, some instances timed out when copying the image to its boot volume.

The GPFS Block Storage service driver now properly detects when the storage back end supports COW mode. You can now spawn multiple VM instances simultaneously.

BZ#2022121

Before this update, the NFS driver blocked attempts to delete OpenStack Storage snapshots in the error state, which prevented users from removing broken snapshot DB entries. With this update, the restriction is removed so that you can clean up failed snapshots.

BZ#2024684

This update fixes a bug that prevented connections to the Ceph storage backend with Ceph client release 15.2.0 (Octopus) and later, affecting Red Hat Ceph Storage 5.0 and later.
A temporary configuration file generated to enable a Ceph connection did not include a '[global]' section marker. This update adds the '[global]' section marker to the temporary file.

The section marker was introduced in Ceph client release 0.94.0 (Hammer). Starting with the Octopus release, Ceph requires the presence of the marker. This fix is backward compatible to Red Hat Ceph Storage 4.x.

BZ#2026290

Before this update, missing parameters caused HAproxy logs to not be forwarded when rsyslog was configured using OpenStack Orchestration. With this update, all required parameters are present, which means that HAProxy logs are included in logs collected and forwarded.

BZ#2027759

This update corrects an error that prevented the proper use of the Block Storage service (cinder) **powermax_port_groups** parameter.

BZ#2036652

Red Hat OpenStack Platform (RHOSP) does not support the use of a fully qualified domain name (FQDN) as the instance display name in a boot server request. The instance display name is passed from the boot server request to the **instance.hostname** field.
A recent update now sanitizes the **instance.hostname** field. The sanitization steps include replacing periods with dashes, a replacement that makes it impossible to continue using the unsupported FQDN instance display names.

This update provides a temporary workaround for customers who use a fully qualified domain name (FQDN) as the instance display name in a boot server request. It limits the scope of the sanitization to cases where the instance display name ends with a period followed by one or more numeric digits.

If you use FQDN as the instance display name in a boot server request, modify your workflow before upgrading to RHOSP 17.

BZ#2038897

This update fixes a bug that caused the **metrics_qdr** service to appear in the HAProxy configuration of an Service Telemetry Framework (STF) deployment on a distributed compute node (DCN) node.

3.5.3. Enhancements

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

BZ#1848200

Snapshots are fully supported in the Shared File Systems service (manila) with the CephFS back ends. Users can create and delete point-in-time backups of their shares through snapshots. Cloud administrators can control the quotas of the number or size of snapshots that users can create.

BZ#1886762

The Block Storage service (cinder) can now use an external NFS share to perform image format conversion of Image service (glance) images on the overcloud Controller nodes. Using this functionality prevents the space on the node from being completely filled during a conversion operation.
See Configuring an external NFS share for conversion .

BZ#1894668

In a DCN deployment, offline volume migration is supported between an edge site and the central site. Volumes may be migrated from an edge site to the central site, and vice versa. However, offline volume migration directly between two edge sites is not supported.

BZ#1921224

This feature allows DCN environments to perform offline volume migration between central and edge sites.

BZ#1949675

This enhancement enables the experimental **rsyslog reopenOnTruncate** setting to ensure that rsyslog immediately recognizes when a logrotation happens on a file. The setting affects every service configured to work with rsyslog.
When **rsyslog reopenOnTruncate** is disabled, rsyslog waits for a log file to fill to its original capacity before consuming any additional logs.

BZ#1969411

Before this update, the Ceph Dashboard could be exposed only on the provisioning network or on a dedicated composable network. Now, Red Hat OpenStack Platform director processes the **CephDashboardNetwork** parameter to define the frontend network used by operators, and limits access to the Ceph Dashboard and monitoring stack to this network.

BZ#1971545

With this update, you can set QoS maximum bandwidth limit, egress direction rules on hardware-offloaded ports in a ML2/OVS deployment. To set the policy, use the normal QoS policy/rules methods.
The back end uses **ip link** commands to enforce the policy instead of the normal OVS QoS engine, because the OVS **meter** action cannot be offloaded. See meter action is not offloaded.

BZ#1977392

With this update, you can use the **openstack undercloud backup** command with the **--db-only** option to create a backup of the database that runs on the undercloud node. You can use that backup to recover the state of the database in the event that it becomes corrupted during the upgrade process.

BZ#1984875

Sometimes a leapp actor blocks the leapp process during an upgrade. You can now remove actors with the new, role-specific **LeappActorsToRemove** parameter.

BZ#1999324

This enhancement updates the default value of the parameter **NovaLiveMigrationPermitAutoConverge** to enabled. When the parameter is enabled, the instance CPU is slowed down until the memory copy process is faster than the instance's memory writes when the migration performance is slow and might not complete.

BZ#1999725

You can now deploy the CephFS NFS gateway (ganesha) on the external network instead of the default dedicated StorageNFS network.

BZ#2029943

This release includes an update of the Block Storage service (cinder) driver for HPE storage arrays. With the updated driver you can use the iSCSI protocol with HPE's Primera products for Primera version 4.2 and later.

BZ#2050154

Red Hat OpenStack Platform (RHOSP) now supports the correct method of updating OVN. For more information, see Optional: Updating the ovn-controller container on all overcloud servers.

3.5.4. Technology previews

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#1952060

With this technology preview update, you can set the following parameters to configure OVS PMD automatic load balance:

- **OvsPmdAutoLb**: Enable/disable the OVS DPDK PMD Auto Load Balance feature. Values: true or false. OVS DPDK uses the default value of false.
- **OvsPmdLoadThreshold**: Set the minimum PMD thread load threshold for OVS DPDK PMD Auto Load Balance feature. Set a value from 0 to 100 to specify the minimum PMD thread load threshold (% of used cycles) of any non-isolated PMD threads when a PMD Auto Load

Balance might be triggered.

- **OvsPmdImprovementThreshold**: Set PMD load variance improvement threshold for OVS DPDK PMD Auto Load Balance feature. Set a value from 0 to 100 to specify the minimum evaluated percentage improvement in load distribution across the non-isolated PMD threads that allows a PMD Auto Load Balance to occur.

- **OvsPmdRebalInterval**: Set PMD auto load balancing interval for OVS DPDK PMD Auto Load Balance feature. Set a value from 0 to 20,000 to specify the minimum time (in minutes) between 2 consecutive PMD Auto Load Balancing iterations.

3.5.5. 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#1978286

Starting with Red Hat Ceph Storage 4, you can enable encryption for all traffic generated by the Ceph daemons over the network.
The secure mode setting for messenger v2 encrypts the communication between Ceph daemons and Ceph clients, effecting an end-to-end encryption.

A new tripleo-heat-templates parameter can enable the on-wire encryption between daemons and clients. To configure Ceph to enable the on-wire encryption between daemons and clients, add the following lines to the overcloud deployment environment files:

```
parameter_defaults:
  CephMsgrSecureMode: true
```

BZ#1982489

With RHOSP 16.2.2, you can configure your undercloud to support both PXE and iPXE boot modes when your architecture includes both POWER (ppc64le) and x86_64 UEFI nodes. For more information, see Configuring a multiple CPU architecture overcloud .

BZ#1984555

This update adds the **CollectdContainerAdditionalCapAdd** variable to the deployment tool. This variable is a comma separated list of additional collectd container capabilities. You can use it to add capabilities to the collectd container.

3.5.6. Known issues

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

BZ#2027787

Disable the advanced-virt-for-rhel-8 repository before you install Red Hat OpenStack Platform (RHOSP) 16.2, update from RHOSP 16.2 to a newer maintenance release, or upgrade from 16.1 to 16.2.
RHOSP hosts do not require the advanced-virt-for-rhel-8 repository. If you do not disable it, dependency issues cause the installation, update, or upgrade to fail. The dependency failures happen because the advanced-virt-for-rhel-8-x86_64-rpms repository is based on RHEL 8.latest, which does not work with RHEL 8.4.

As a workaround, disable the repositories. Perform the steps appropriate for your installation, update, or upgrade scenario.

- Scenario: new 16.2 installation or update from 16.2 to later version of 16.2.
  $ subscription-manager repos --disable advanced-virt-for-rhel-8-x86_64-rpms

  $ dnf module disable -y virt:av

  $ dnf module enable -y virt:rhel

- Scenario: upgrade from 16.1→16.2.
  $ subscription-manager repos --disable advanced-virt-for-rhel-8-x86_64-rpms

  $ dnf module disable -y virt:8.2

  $ dnf module enable -y virt:rhel

BZ#2050765

There is currently a known issue in which the OpenStack **tripleo validator run** command produces errors. In the command output the Status_by_Host column contains the value, 'No host matched'. This error occurs with all **openstack validator run** calls because no ansible inventory was available for the undercloud and is not limited to a single validator group. This error also impacts overcloud validations. The root cause of this is a regression in the tripleo client code where tripleo-ansible-inventory is no longer called in real time.
As a result, in updates to Red Hat OpenStack Platform (RHOSP) 16.2, the **tripleo validator run --group pre-introspection** command fails on all tests.

Workaround: create a file called inventory.yaml by manually running **tripleo-ansible-inventory --static-yaml-inventory inventory.yaml**. Then, run the validation using the **-i inventory.yaml** argument:

```
# tripleo-ansible-inventory --static-yaml-inventory inventory.yaml
# openstack tripleo validator run  --group pre-introspection -i inventory.yaml
```

For more information, see the Red Hat Knowledgebase solution openstack tripleo validator run command produces error.

3.5.7. Deprecated functionality

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

BZ#2023517

The collectd plugin write_redis has been deprecated in Red Hat OpenStack Platform (RHOSP) 16.2 and will be removed in RHOSP 17.0.

3.6. RED HAT OPENSTACK PLATFORM 16.2.1 MAINTENANCE RELEASE - DECEMBER 09, 2021

Consider the following updates in Red Hat OpenStack Platform (RHOSP) 16.2.1 when you deploy this RHOSP release.

3.6.1. Advisory list

This release of Red Hat OpenStack Platform (RHOSP) includes the following advisories:

RHBA-2021:5067

Release of components for OSP 16.2

RHEA-2021:5068

Red Hat OpenStack Platform 16.2.1 director images bug fix advisory

RHEA-2021:5069

Red Hat OpenStack Platform 16.2 containers bug fix advisory

3.6.2. Bug fixes

These bugs were fixed in this release of Red Hat OpenStack Platform:

BZ#1977442

Before this update, the ML2/OVN container started, but the service did not work correctly. The cause for this problem was that the ML2/OVN controller container did not include a configuration to support TLS. With this update, the ML2/OVN container configuration has been updated to include all of the correct configuration, and it now works correctly.

BZ#2003708

This update fixes a bug that prevented migration from ML2/OVS with VXLAN to ML2/OVN with Geneve in RHOSP 16.2. You can now migrate from ML2/OVS with VXLAN to ML2/OVN with Geneve in RHOSP 16.2.

BZ#2005404

Before this update, the **certmonger** package was dropped from the minimal image, which made it impossible to deploy Red Hat Ceph Storage nodes.
With this update, the **certmonger** package has been added back to the image, and Red Hat Ceph Storage nodes can now be deployed.

BZ#2007268

Before this update, a lock handling issue prevented IPMI-based nodes from recording the hardware vendor as part of power state synchronization. This issue caused the power state synchronization to fail, and nodes that used the **ipmi** hardware type entered the **Maintenance** state. With this update, the lock is handled correctly and the power state synchronization for bare metal nodes that use the **ipmi** hardware type work correctly and no locking errors occur.

BZ#2008981

Before this update, removal of the **python2** packages for the Red Hat Enterprise Linux (RHEL) in-place upgrade tool, LEAPP, was unsuccessful. This failure was caused by a DNF **exclude** option that retained the LEAPP packages. With this update, automation has now been included to ensure that the necessary LEAPP packages are successfully removed.

BZ#2019178

Before this update, an upgradable **mariadb-server** package in the RHEL repository caused the package manager to upgrade the **mariadb-server** package on the host, interfering with the containerized **mariadb-server** that pre-exists on the same host. With this update, the Red Hat OpenStack Platform (RHOSP) director removes the **mariadb-server** package from any hosts which also have the containerized MariaDB, and the RHOSP FFU process continues.

3.6.3. 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#1989820

When configuring bandwidth-aware scheduling of SR-IOV workloads, use the heat parameter, **resource_provider_hypervisors**. This parameter defines a paired list **<network_device>:<hypervisor>**. Ensure that you use a fully qualified domain name (FQDN) to define each hypervisor.

BZ#2007255

With this update, the memory limit for the **collectd** container has been increased to 512 MB. When this limit is exceeded, the container is restarted.

3.6.4. Technology previews

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#1892796

In Red Hat OpenStack Platform (RHOSP) 16.2.1, a technology preview is available that supports Intel Columbiaville E810 NICs in NFV deployments, with the following recommendations:

- You cannot configure Dynamic Device Personalization (DDP) in the heat template.
- Live migration is unsupported.
- Virtual Function (VF) rate limiting is unsupported in RHOSP 16.

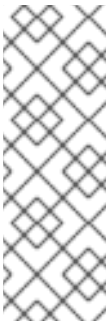

NOTE

In Red Hat OpenStack Platform (RHOSP) 16.2.3, Intel Columbiaville E810 NICs in NFV deployments are now fully supported with the following known issues:

- OVS is only able to use one VF. (BZ#1978622)
- PMD cores running in kernel space for OVN-DPDK. (BZ#2049442)

3.6.5. Known issues

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

BZ#1966157

There is a limitation when using ML2/OVN with **network_type geneve** with a Mellanox adapter on a Compute node that has more than one instance on the geneve network. The floating IP of only one of the instances is reachable.

BZ#2003708

RHOSP does not yet support ML2/OVN with VXLAN networks. The migration process includes steps to convert VXLAN networks to Geneve. When the migration target version is RHOSP 16.2.0, a bug prevents the expected VXLAN to Geneve conversion, and the networks remain configured as VXLAN.

This bug affects only migrations to ML2/OVN on RHOSP 16.2. It does not affect migrations to ML2/OVN on RHOSP 16.1.

3.7. RED HAT OPENSTACK PLATFORM 16.2 GA

Consider the following updates in Red Hat OpenStack Platform (RHOSP) 16.2.6 when you deploy this RHOSP release.

3.7.1. Advisory list

This release of Red Hat OpenStack Platform (RHOSP) includes the following advisories:

RHEA-2021:3483

Release of components for OSP 16.2

RHEA-2021:3485

Red Hat OpenStack Platform 16.2 deployment images

RHEA-2021:3486

Release of containers for OSP 16.2 director operator tech preview

RHSA-2021:3487

Moderate: Red Hat OpenStack Platform 16.2 (etcd) security update

RHSA-2021:3488

Important: Red Hat OpenStack Platform 16.2 (openstack-neutron) security update

RHEA-2021:3489

Release of components for OSP 16.2 - Containers

RHSA-2021:3490

Moderate: Red Hat OpenStack Platform 16.2 (python-django20) security update

3.7.2. Bug fixes

These bugs were fixed in this release of Red Hat OpenStack Platform:

BZ#1690726

Before this update, writing an image to RBD could be very slow. This update improves the process for writing an image to RBD, which improves the time it takes for images to be written to RBD.

BZ#1772531

Typically, when you create an encrypted volume from a snapshot of an encrypted volume, the source volume is the same size or smaller than the destination volume.
In previous releases, if you created an encrypted volume from a snapshot of an encrypted volume, and the destination volume was close to or equal to the size of the source volume, the Block Storage service (cinder) silently truncated the data in the new destination volume.

With this release, the Block Storage service calculates the size of the destination volume to include the current size of the encryption header, which eliminates the data truncation.

BZ#1844372

Before this update, when you resized or migrated an instance that had a vGPU flavor you needed to rebuild the instance manually to re-allocate the vGPU resources. With this update, instances with a vGPU flavor are automatically re-allocated the vGPU resources after resize and cold migration

operations.

BZ#1849843

Previously, the Shared File Systems service (manila) API that brings external shares into service management did not check for duplicated export locations. An existing share brought into the service multiple times results in an inconsistent state.
With this release, the API evaluates the export locations of known or existing shares before allowing external shares to be managed, and prevents existing shares from being erroneously brought into the Shared File Systems service again.

BZ#1851051

Before this update, RBD performance was degraded when multiple instances were launched simultaneously. This was due to the Image service (glance) starting multiple threads to perform the same copying operation. This update resolves the issue.

BZ#1851797

This update fixes an Image service (glance) configuration error that prevented users from creating a virtual machine with watchdog by setting flavor metadata.

BZ#1884322

In prior releases, you could not delete snapshots that have snapshot dependencies. With this release, you can delete snapshots that have snapshot dependencies.

BZ#1888105

When multiple storage back ends are configured on the Shared File Systems service (manila), each storage back end might support a different storage protocol. Before this update, the Shared File Systems service scheduler did not consider the storage protocol and capability of the shared storage back ends when deciding where to place them, which caused share provisioning to fail. With this update, the Shared File Systems service scheduler now automatically considers the share type extra specs with the storage protocol, which makes it possible to use the **CapabilitiesFilter** to compare storage back-end capabilities and provision shares successfully.

BZ#1910508

Before this update, validation results were not logged and validation artifacts were not collected because the permissions required to access the requested logging directory were not granted. This update resolves the issue, and validation results are successfully logged and validation artifacts are collected.

BZ#1913671

The Unisphere for PowerMax REST endpoints have changed from 91 to 92. This update changes how URIs are created, to allow for full coverage of all possible Unisphere REST API endpoints.

BZ#1919855

When an instance is created, the Compute service (nova) sanitizes the instance display name to generate a valid hostname when DNS integration is enabled in the Networking service (neutron). Before this update, the sanitization did not replace periods ('.') in instance names, for example, 'rhel-8.4'. This could result in display names being recognized as Fully Qualified Domain Names (FQDNs) which produced invalid hostnames. When instance names contained periods and DNS integration was enabled in the Networking service, the Networking service rejected the invalid hostname, which resulted in a failure to create the instance and a HTTP 500 server error from the Compute service.

With this update, periods are now replaced by hyphens in instance names to prevent hostnames being parsed as FQDNs. You can continue to use free-form strings for instance display names.

BZ#1923975

Before this update, some exceptions were not caught during connections to iSCSI portals, such as failures in **iscsiadm -m session**. This caused **_connect_vol** threads to abort unexpectedly in some

failure patterns, which caused subsequent steps to hang while waiting for results from **_connect_vol** threads. This update ensures that any exceptions during connections to iSCSI portals are handled correctly in the **_connect_vol** method, to avoid unhandled exceptions during connecting to iSCSI portals, and unexpected aborts that have no updated thread results.

BZ#1935154

This update adds Challenge Handshake Authentication Protocol (CHAP) support to the Dell EMC PowerStore driver. PowerStore can now be used with enabled CHAP as a storage back end.

BZ#1939394

Before this update, the NetApp SolidFire driver created a duplicate volume when the API response is lost due to a connection error and the driver retries the API request. This occurred when the SolidFire back end successfully received and processed a create volume operation, but failed to deliver the response back to the driver. This update resolves the issue by:

1. Checking if the volume name already exists in the back end before trying to create it. If a volume is found, an exception is raised and the process is aborted.

2. Checking for volume creation right after a read timeout is detected, to prevent invalid API calls.

3. Adding the ´sf_volume_create_timeout´ option to the SolidFire driver, to allow users to set the appropriate timeout value for their environment.

BZ#1942531

Before this update, execution of the validation package **check-latest-packages-version** was slow. This update resolves the issue.

BZ#1942717

This release supports port filtering for the Dell EMC XtremeIO driver for the Block Storage service (cinder).

BZ#1953749

Before this update, if PowerStore ports were configured for multiple purposes, such as for iSCSI or Replication, the driver reported that it could find no accessible iSCSI targets. This was because the REST filter was wrong. This update fixes the PowerStore iSCSI targets filtering.

BZ#1956370

Before this update, when iSCSI or FC targets, such as ESXi, were not connected to the RHOSP host, the attach volume operation waited until it timed out. With this update, a new option to support ports filtering has been added to the Dell EMC XtremIO driver for the Block Storage service (cinder).

BZ#1959853

Before this update, the validation variable in one code path was referenced but never assigned, which resulted in an unhandled exception during validation. This has been fixed.

BZ#1960185

Before this update, Ansible redirected output to all registered non-stdout callback plug-ins by default, which resulted in VF callbacks processing information from other processes that used **ansible runtime**. This issue has been resolved and the output of other processes is no longer stored in the validations logging directory.

BZ#1972774

This update fixes an issue that caused Networking service (neutron) agents, such as Networking service DHCP, to fail when they tried to create resources in OVN because ML2/OVN prevented RPC workers from connecting the OVN southbound database.

BZ#1974979

Before this update, the Shared File Systems service (manila) dashboard had dynamic form elements whose names could potentially cause the forms to become unresponsive. This meant that the creation of share groups, share networks, and shares within share networks did not work.
With this update, dynamic elements whose names could be problematic are encoded, which means that creation of share groups, share networks, and shares within share networks functions normally.

BZ#1976693

The Shared File Systems service (manila) uses the CephFS volume client to communicate with Ceph Storage clusters. Previously, the CephFS volume client package aborted while creating or deleting file systems.
The aborted operations caused the manila-share process within the Shared File Systems service to restart, which caused shares that were being provisioned or deleted to be stuck in **creating** or **deleting** states, respectively.

With this release, the CephFS volume client package no longer aborts provisioning or deletion requests, and the manila-share process does not restart during these operations.

BZ#1978158

This update fixes an issue that caused Networking service (neutron) agents, such as Networking service DHCP, to fail when they tried to create resources in OVN. This was caused by residual data left in the OVN databases when QoS rules were created for floating IPs.
This update eliminates the residual data and fixes the problem.

BZ#1985717

This update fixes a known issue where the Open Virtual Network (OVN) Metadata service was not available to VM instances bound to an SR-IOV virtual function. The issue did not affect network function but these instances did not receive their SSH keys in the absence of a Metadata service connection.
The metadata service connectivity for SR-IOV ports now functions correctly.

BZ#1987092

Before this fix, grub2 tooling wrote kernel argument changes to /boot/grub2/grubenv. This file was not available to UEFI boot systems, and caused kernel argument changes not to persist across reboots on UEFI boot nodes.
This fix changes both the /boot/grub2/grubenv file and the /boot/efi/EFI/redhat/grubenv files when you make kernel argument changes. As a result, RHOSP director now applies persistent Kernel argument changes for UEFI boot nodes.

3.7.3. Enhancements

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

BZ#1714772

This enhancement adds support for Entrust nShield HSM deployment in high availability mode with OpenStack Key Manager (barbican).

BZ#1866741

Images with null bytes take up a lot of space. With this release, you can enable sparse image upload and save space when you upload images. Sparse image upload is supported only with Ceph RBD.

BZ#1868940

This enhancement adds the type **HostDomain**. **HostDomain** is the same as **HostAddress** with the

added support of the underscore character - RFC1033. Systems such as DomainKeys and service records use the underscore. The Compute service can use the **HostDomain** type to define **live_migration_inbound_addr**.

BZ#1880141

Red Hat OpenStack Platform 16.2 includes support for Single Root Input/Output Virtualization (SR-IOV) and Data Plane Development Kit (DPDK) workloads on AMD hosts.

BZ#1897890

This enhancement improves the efficiency, performance, and execution time of deployment and update tasks for environments with a large number of roles. The logging output of the deployment process has been improved to include task IDs for better tracking of specific task executions, which can occur at different times. You can use the task IDs to correlate timing and execution when you troubleshoot executions.

BZ#1900723

During stack update the **KernelArgs** could be modified or appended. You must perform a reboot of the affected nodes manually.
For example, if the current deployment has the following configuration, it is possible to change **hugepages=64**, or add or remove arguments during the stack update:

```
`KernelArgs: "default_hugepagesz=1GB hugepagesz=1G hugepages=32 intel_iommu=on iommu=pt isolcpus=1-11,13-23"
```

For example:

```
KernelArgs: "default_hugepagesz=1GB hugepagesz=1G hugepages=64 intel_iommu=on iommu=pt isolcpus=1-24"
KernelArgs: "isolcpus=1-11,13-23"
```

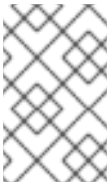

NOTE

Complete removal of **KernelArgs** during the update is not supported. Also **KernelArgs** could be newly added as well to an existing overcloud node, however the reboot was triggered in this case.

BZ#1920229

With this enhancement, you can improve the performance of live migrations by using the following new parameters:

- **NovaLiveMigrationPermitPostCopy** - When enabled, the instance is activated on the destination node before migration is complete, and an upper bound is set on the memory that needs to be transferred, which improves the live migration of larger instances. This parameter is enabled by default.
- **NovaLiveMigrationPermitAutoConverge** - When enabled, if an on-going live migration is progressing slowly the instance CPU is throttled until the memory copy process is faster than the instance's memory writes. This parameter is disabled by default. To enable **NovaLiveMigrationPermitAutoConverge**, add the following configuration to an environment file:

```
parameter_defaults:
  ComputeParameters:
    NovaLiveMigrationPermitAutoConverge: true
```

BZ#1926721

This enhancement improves the performance and application of the **check-latest-packages-version** validation.

BZ#1926725

This enhancement adds new validation for **tripleo-latest-packages-version**. This validation checks if the listed **tripleo** packages are up to date with repositories.

BZ#1938568

Before this update, the **PluginInstanceFormat** parameter for **collectd** could specify only one of the following values: 'none', 'name', 'uuid', or 'metadata'. After this update, the **PluginInstanceFormat** parameter for **collectd** can now specify more than one value, which results in more information being sent in the **plugin_instance** label of **collectd** metrics.

BZ#1977392

With this update, you can use the **openstack undercloud backup** command with the **--db-only** option to create a backup of the database that runs on the undercloud node. You can use that backup to recover the state of the database in the event that it becomes corrupted during the upgrade process.

3.7.4. Technology previews

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#1825895

In Red Hat OpenStack Platform 16.2, a technology preview is available that supports Precision Time Protocol (PTP) with Timemaster.

BZ#1925999

The Red Hat OpenStack Platform (RHOSP) director Operator creates a set of custom resource definitions (CRDs) on top of Red Hat OpenShift Container Platform to manage resources normally created by the RHOSP undercloud. CRDs are split into two types for hardware provisioning and software configuration. The operator includes CRDs to create and manage overcloud networks, manage IP addresses, create VM sets for RHOSP Controllers, and create bare metal sets for RHOSP Computes.
For Technology Preview, the software configuration is accomplished with an OpenStackClient pod, which uses traditional RHOSP or TripleO interfaces and CLI commands. Work is ongoing to produce a more scalable Heat-to-Ansible playbook deployment workflow within the RHOSP director Operator.

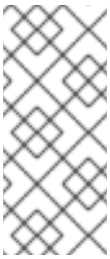

NOTE

The Red Hat OpenStack Platform (RHOSP) director operator became a fully supported feature shortly after the release of Red Hat OpenStack Platform (RHOSP) 16.2.4 Maintenance Release, on December 13, 2022. For deployment information, see RHBA-2022:8952, Release of containers for Red Hat OpenStack Platform 16.2.4 director operator.

3.7.5. 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#1654408

For Image Service (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 the **glance-api.conf** file.
The Image Service must have a staging area when you use the **glance-direct** import method. Set the **node_staging_uri** option in the **DEFAULT** section of the **glance-api.conf** file to **file://<absolute-directory-path>**. This path must be on a shared file system that is available to all Image service API nodes.

BZ#1906028

The **python-networking-fujitsu** package is not included with RHOSP 16.2.

BZ#1961784

In this release, EFI bootloader assets for whole-disk images are preserved during deployment, so the shim bootloader is no longer overwritten. This ensures that Secure Boot is switched on after deployment.

BZ#1978286

Starting with Red Hat Ceph Storage 4, you can enable encryption for all traffic generated by the Ceph daemons over the network.
The secure mode setting for messenger v2 encrypts the communication between Ceph daemons and Ceph clients, effecting an end-to-end encryption.

A new tripleo-heat-templates parameter can enable the on-wire encryption between daemons and clients. To configure Ceph to enable the on-wire encryption between daemons and clients, add the following lines to the overcloud deployment environment files:

```
parameter_defaults:
  CephMsgrSecureMode: true
```

BZ#1989820

When configuring bandwidth-aware scheduling of SR-IOV workloads, use the heat parameter, **resource_provider_hypervisors**. This parameter defines a paired list **<network_device>:<hypervisor>**`. Ensure that you use a fully qualified domain name (FQDN) to define each hypervisor.

BZ#1992655

In previous releases, administrators had to add the **ceph** plugin to the **CollectdExtraPlugins** parameter in their custom environment files.
With this release, the **ceph** plugin loads automatically on Ceph Storage nodes. Therefore, before you upgrade from Red Hat OpenStack Platform 13 to 16.2, you must remove the **ceph** plugin from the **CollectdExtraPlugins** parameter in your custom environment files.

3.7.6. Known issues

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

BZ#2111871

Significant connectivity loss can affect your workloads after certain updates to RHOSP 16.2.2 and 16.2.3.

The bug affects your deployment when you update RHOSP to 16.2.2 or 16.2.3 from 16.2.0, 16.2.1, or from any 16.1 release. To verify if your update path is affected, see the article [1].

What to do

If your planned update path is affected by the bug and you do not have a strong reason to update now, Red Hat recommends that you wait for the planned release of a fix in RHOSP 16.2.4.

If your planned update path is affected by the bug and you must update now to 16.2.2 or 16.2.3, contact your Red Hat support representative to see if your deployment is compatible with a Hot Fix that addresses the bug.

Bug details

The bug is triggered by a database schema change in OVN 21.12, which is introduced in RHOSP 162.2. and 16.2.3. OVN 21.12 contains a new column that is not present in earlier versions. OVN database schema changes should not cause a problem in OpenStack, but this particular change is affected by a bug.

In particular, instance connectivity is lost for a variable amount of time (from 20 seconds to 3 minutes) when you run the following command:

```
$ openstack overcloud external-update run --stack overcloud --tags ovn
```

A fix for this bug is scheduled for RHOSP 16.2.4.

[1] Follow the steps in this article to determine if your update path is affected:

https://access.redhat.com/solutions/6974338

BZ#2129445

Do not update from RHOSP 116.2.0 to 16.2.2 or 16.2.3 until you evaluate your risk of serious impact from a libvirt version incompatibility. To complete this evaluation, check the libvirt package in the **nova_libvirt** container in all compute nodes:

```
$ sudo podman exec nova_libvirt rpm -q libvirt
```

If the libvirt version is 7.0, the deployment IS NOT affected by the bug. You can perform the update.

If the libvirt version is 7.6, the deployment is affected by the buy. Your update is at risk.

If you learn that your update is at risk from the libvirt incompatibility, choose one of these options:

Wait: Update directly to RHOSP 16.2.4 when it is released. It includes a fix for the incompatibility issue. This is the preferred option if you can postpone the update. In 16.2.4 you can perform the manual steps outlined in KCS [1] without any hot fix, because the option skip_hypervisor_version_check_on_lm is included in 16.2.4.

Hot fix: Contact your Red Hat support representative to explore whether your environment is compatible for a hot fix patch that resolves the issue. Use this option if there is a strong business need for an immediate update.

If you already performed the update with the version incompatibility, see the KCS article [2] for guidance on fixing the problem.

[1] Article helps you verify if your update is at risk (same step as above). Also includes steps you can perform to update to 16.2.1, 16.2.2, or 16.2.3 if you choose the Hot Fix option. See https://access.redhat.com/solutions/6972451

[2] Article helps you fix your deployment if you already completed an update to 16.2.1, 16.2.2, or 16.2.3 and your deployment is affected by the incompatibility. See https://access.redhat.com/solutions/6969430

BZ#1975240

Starting with Red Hat Enterprise Linux (RHEL) version 8.3, support for the Intel Transactional Synchronization Extensions (TSX) feature is disabled by default. Currently, this causes instance live migration to fail when migrating from hosts where the TSX kernel argument is enabled to hosts where the TSX kernel argument is disabled.
This impact applies only to Intel hosts that support the TSX feature. For more information about the CPUs that are affected by this issue, see Affected Configurations.

For more information, see the Red Hat Knowledgebase solution Guidance on Intel TSX impact on OpenStack guests.

BZ#1983748

In Red Hat OpenStack Platform (RHOSP) deployments that use the Modular Layer 2 plug-in with the Open vSwitch (ML2/OVS) mechanism driver, there is currently a known issue where the Orchestration service (heat) parameter, **NeutronL3AgentAvailabilityZone** does not set the relevant Neutron L3 agent parameter correctly.
Workaround: use a custom hieradata statement to set this value. In the example that follows, replace **[ROLE]** with the composable role name that is appropriate for your site.

Example

```
[ROLE]ExtraConfig:
  neutron::agents::l3::availability_zone: role_availability_zone
```

For more information, see Puppet: Customizing hieradata for roles in the *Advanced Overcloud Customization* guide.

BZ#1986423

Rebooting a node with a virtual function (VF) attached to OVS-DPDK (vfio-pci driver) results in VF uninitialized on that physical function (PF). As a result, virtual machines are unable to use the VFs from that PF. If a second VF is used for another OSP network, it will not function as expected after reboot.
Workaround: perform the following steps on the Compute node before you reboot the node:

1. Delete the file "/etc/udev/rules.d/70-os-net-config-sriov.rules"
2. Modify the "Before" criteria of "/etc/systemd/system/sriov_config.service" file to add "network-pre.target" The modified "Before" should look like:

   ```
   Before=network-pre.target openvswitch.service
   ```

   The workaround fixes the issue and all the VFs initialize correctly.

BZ#2003708

RHOSP does not yet support ML2/OVN with VXLAN networks. The migration process includes steps

to convert VXLAN networks to Geneve. When the migration target version is RHOSP 16.2.0, a bug prevents the expected VXLAN to Geneve conversion, and the networks remain configured as VXLAN.
This bug affects only migrations to ML2/OVN on RHOSP 16.2. It does not affect migrations to ML2/OVN on RHOSP 16.1.

BZ#1855423

This update fixes a bug that prevented fast forward upgrades (FFU) of instance HA environments from RHOSP 13 to RHOSP 16.1.

BZ#1856901

There are known limitations for Mellanox ConnectX-5 adapter cards in VF link aggregation group (LAG) mode in OVS OFFLOAD deployments, SRIOV Switchdev mode.
When at least one VF of any physical function (PF) is still bound or attached to a VM, an internal firmware error occurs when attempting to disable single-root input/output virtualization (SR-IOV), and when unbinding PF using a function such as **ifdown** and **ip link**.

Workaround: Reboot the node to restore the bond. For more information, see the Red Hat Knowledgebase solution Mellanox ConnectX-5 internal error when removing PF from the bond or disabling SR-IOV.

VF LAG mode with OVS OFFLOAD, SRIOV switchdev mode is not supported, if you use the **mstconfig** tool to set a value higher than 64 for the **NUM_OF_VFS** parameter in the Firmware configuration. Currently, there is no workaround available.

BZ#2109597

There is a hardware (HW) limitation with CX-5. Every network traffic flow has a direction in HW, either transmit (TX) or receive (RX). If the source port of the flow is a virtual function (VF), then it is also a TX flow in HW. CX-5 cannot pop VLAN on the TX path, which prevents offloading the flow with **pop_vlan** to the HW.

3.7.7. Deprecated functionality

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

BZ#1868673

For Distributed Compute Node deployments which use storage, **dcn-hci.yaml** has been renamed to **dcn-storage.yaml** because DCN sites with storage have the option of not using HCI (Hyper-Converged Infrastructure). **dcn-hci.yaml** is deprecated but will remain in the environments directory for backwards compatibility. **dcn-hci.yaml** will be removed in Red Hat OpenStack platform 17. **dcn-storage.yaml** should be used in place of **dcn-hci.yaml**.

BZ#1984484

Block Storage service (cinder) backup with Google Cloud Storage is being deprecated. Support will be removed in the next major release.

BZ#1984887

In this release, Block Storage service (cinder) backup support for Google Cloud Services (GCS) has been deprecated. Support will be removed in Red Hat OpenStack Platform (RHOSP) 17.0.

BZ#1990802

In Red Hat Openstack Platform (RHOSP) 16.2, support for the QXL video model is deprecated, due to the removal of support for the Spice graphics software in RHEL 9. This will cause issues for instances that use QXL when migrating from RHEL-8 to RHEL-9. Red Hat recommends using the

virtio video model for both UEFI and BIOS instances, instead of **qxl**. When creating a new instance from an image, set the video model before launching the new instance:

```
$ openstack image set --property hw_video_model=virtio <image>
```

To update the video model for existing instances that use the QXL video model:

1. Stop the instance.
2. Snapshot the instance.
3. Update the image metadata on the instance snapshot image to include the property **hw_video_model=virtio**.
4. Create a new instance using the instance snapshot.
   For more information on supported video models, see Image configuration parameters in the *Creating and Managing Images* guide.

CHAPTER 4. TECHNICAL NOTES

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

4.1. RHBA-2023:1763 – RELEASE OF COMPONENTS FOR OSP 16.2.5

Changes to the openstack-aodh component:

- Before this update, the Alarming service (aodh) used a deprecated Gnocchi API to aggregate metrics, which sometimes caused Gnocchi to display incorrect CPU use values. With this update, Gnocchi displays the correct metrics by performing calculation, transformation, and aggregation of metrics dynamically. (BZ#2133030)

Changes to the openstack-cinder component:

- Before this update, the Block Storage service PowerMax driver failed to properly handle volume snapshots based on the PowerMax legacy snapshot identification method of a generation number. This caused the failure of attempts to manage a snapshot identified by a generation number. With this update, the PowerMax driver properly handles and manages snapshots identified by a generation number. (BZ#2172897)

- Before this update, the **pure_iscsi_cidr** parameter of the Block Storage service (cinder) PureISCSIDriver did not support IPv6 addresses. With this update, the **pure_iscsi_cidr** parameter of the Block Storage service PureISCSIDriver supports IPv6 addresses and this driver provides a new parameter called **pure_iscsi_cidr_list**, which supports a list of networks. (BZ#2142282) Changes to the openstack-heat component:

- Before this update, a change in the **auth_encryption_key** parameter caused an inability to delete existing Heat stacks. With this update, Heat allows for changes in the **auth_encryption_key** parameter when deleting existing Heat stacks. Heat ignores objects that cannot be decrypted when deleting Heat stacks. (BZ#2142684)

Changes to the openstack-octavia component:

- Before this update, inadequate TCP buffer sizes resulted in out of memory warnings for TCP in amphora. The smaller TCP buffer size had a potential negative impact on TCP flows with large payloads. This update increases the size of the TCP buffers in amphora, improving the reliability of the TCP connections. This resolves the issue. (BZ#2151893)

- Before this update, a race condition occurred in Octavia that might have caused load balancers that use the OVN provider to become stuck in PENDING DELETE under certain conditions. This caused the load balancer to be immutable and unable to update. With this update, the race condition is fixed to resolve the issue. (BZ#2165032)

Changes to the openstack-tripleo-heat-templates component:

- Before this update, the multipath daemon running in a container did not detect changes in the underlying multipath devices on the host. This resulted in failure of Block Storage operations such as resizing an online volume. With this update, the container running the multipath daemon is kept synchronized with multipath devices on the host so that Block Storage operations on multipath volumes function correctly. (BZ#2165494)

- With this update, operators have the ability to configure the number of metadata agent workers using the **NeutronMetadataWorkers** parameter in the THT. Each OVN metadata agent worker creates a connection to the OVN southbound database. For optimal scaling, it is good practice

to set the worker count to 1 to avoid overloading the database. (BZ#2154361)

- Before this update, an issue caused glance_api cron jobs from being triggered. With this update, the issue is resolved. (BZ#2155987)

Changes to the puppet-tripleo component:

- Before this update, customized Heat policy rules were not applied to the **heat-engine** service. This omission caused some of the customized Heat policy rules defined by the **HeatApiPolicies** paramter to be ignored. With this update, director now generates the customized policy file for all Heat services, including **heat-api**, **heat-api-cfn**, and **heat-engine**. All customized Heat policy rules are now applied. This resolves the issue. (BZ#2113819)

Changes to the tripleo-ansible component:

- Before this update, deploying Block Storage multipath with a custom configuration containing a "blacklist_exceptions" section failed due to an error in the multipath configuration. With this update, custom multipath configurations are handled correctly and therefore Block Storage multipath can be deployed with a custom configuration containing a "blacklist_exceptions" section. (BZ#2159555)

- Before this update, endpoints of disabled telemetry services were not cleaned up after an upgrade. This omission did not impact the cloud. With this update, upgrades delete obsolete telemetry endpoints. (BZ#1876045) Changes to the documentation component:

- The technology preview support added in RHOSP 16.1 for configuring NVDIMM Compute nodes to provide persistent memory for instances has been deprecated in RHOSP 16.2.5, and will be removed in RHOSP 17.0. Red Hat is removing support for persistent memory from RHOSP 17.0 and future releases in response to the announcement by the Intel Corporation on July 28, 2022 that they are discontinuing investment in their Intel® Optane™ business:

  - Intel® Optane™ Business Update: What Does This Mean for Warranty and Support

  - Intel® Product Change Notification #119311-00

  Cloud operators must ensure that no instances use the vPMEM feature before upgrading to 17.1. (BZ#2187380)

4.2. RHBA-2022:8794 – RELEASE OF COMPONENTS FOR OSP 16.2.4

Changes to the openstack-cinder component:

- Before this update, a race condition occurred when the Compute service (nova) requested that the Block Storage service (cinder) detach a volume and then an external request was made to delete this same volume. This could result in the volume being deleted first before it was detached, which prevented the Compute service from removing this non-existent volume. With this update, if the Compute service requests that the Block Storage service detach a volume and then an external request is made to delete this same volume, this volume is always detached first and then it is deleted. (BZ#2057002)

- This update fixes the following PowerMax Block Storage (cinder) driver issues:

  - Before this update, the PowerMax Block Storage driver deleted all non-temporary snapshots during the **do_sync_check** operation. This update adds a check to determine if a snapshot must be deleted. This ensures that the **do_sync_check** operation does not indiscriminately delete non-temporary snapshots.

    - Before this update, the PowerMax Block Storage driver used case-sensitive conditions which could return errors when modifying the storage group. With this update, these conditions are case-insensitive and storage groups can be modified successfully. (BZ#2126616)

- This update fixes a bug that deleted the existing Block Storage (cinder) backup record when a backup record was imported for an existing **backup_id**. (BZ#2131386)

Changes to the openstack-ironic-python-agent component:

- Deploying RHEL 8.6 images in UEFI mode failed when using the ironic-python-agent, because the ironic-python-agent service did not understand the RHEL 8.6 UEFI boot loader hint file. With this update, you can now deploy RHEL 8.6 in UEFI mode. (BZ#2134529)

Changes to the openstack-nova component:

- This RHOSP 16.2.4 update makes it possible for you to correct a libvirt version incompatibility before updating to RHOSP 16.2.4. If you have the incompatibility issue and do not addresss it before updating to RHOSP 16.2.4, the update might leave instances in an unmanageable state. Before updating to 16.2.4, see the KCS article Workaround for a libvirt version-compat issue (bug 2109350) when updating RHOSP 16.2.0.

  Perform the steps in the article to determine whether your update path is affected by the libvirt incompatibility issue. If it is affected, perform the steps to resolve the issue. (BZ#2109350)

Changes to the openstack-octavia component:

- Before this update, VM instances (amphorae) for the Red Hat OpenStack Platform (RHOSP) Load-balancing service (octavia) could experience performance issues when a lot of connections filled the network connection tracking (conntrack) table. The cause for this was that conntrack was enabled for all packet types, including TCP, which does not require conntrack. In RHOSP 16.2.4, amphora performance has improved, because conntrack is disabled for TCP packets and is only enabled for UDP and SCTP packets. (BZ#2123226)

- Before this update, an SELinux issue triggered errors with Red Hat OpenStack Platform (RHOSP) Load-balancing service (octavia) ICMP health monitors that used the amphora provider driver. In RHOSP 16.2.4, this issue has been fixed and ICMP health monitors function properly. (BZ#2123318)

Changes to the openstack-tripleo-common component:

- RHSA-2022:6969 introduced the process to clean up files in the /var/lib/mistral directory in the undercloud but the process consistently failed when the Load-balancing service (octavia) or Red Hat Ceph Storage was enabled because these services created additional directories, which the cleanup process could not properly remove. Some deployment actions, such as scale out, consistently failed if the Load-balancing service or Ceph Storage was enabled. With this update, Mistral no longer executes the cleanup. Users must manually delete files if they want to enforce the reduced permission of the files in the /var/lib/mistral directory. Deployment actions no longer fail because of a permission error. (BZ#2137484)

- This enhancement adds a method for pulling down third-party containers by introducing a Jinja template processing approach and adding a **template basedir** parameter, which is required by the Jinja loader, to the BaseImageManager. With this update, pulling down the Ceph-related containers is now optional. You can avoid pulling down the Ceph-related containers by setting the the **ceph_images** Boolean value to **False**. (BZ#1933751)

- Before this update, when the number of objects in the Object Storage service (swift) container

for the overcloud exceeded 10,000, only the first 10,000 objects in the config-download directory were cleaned up during a delete operation, and the remaining objects prevented the container from getting deleted. With this update, there is added handling for cases where there are more than 10,000 objects in the Object Storage service container for the overcloud. (BZ#2119145)

Changes to the openstack-tripleo-heat-templates component:

- If you are prepared to take your data plane offline, you can now upgrade the whole overcloud at once. With this enhancement, you complete the upgrade much faster, at the cost of some data plane downtime. For more information, see Speeding up an overcloud upgrade . (BZ#2027851)

- This enhancement adds new configuration options for the Networking service (neutron) logging service plug-in. With this update, there is added support for network security group logging, and you can configure Networking service logging by using the following new parameters in RHOSP Orchestration service (heat) templates:

- Options for Layer 3 (L3) agents:
  - **NeutronL3AgentLoggingRateLimit**
  - **NeutronL3AgentLoggingBurstLimit**
  - **NeutronL3AgentLoggingLocalOutputLogBase**
- Options for for Open vSwitch (OVS) agents:
  - **NeutronOVSAgentLoggingRateLimit**
  - **NeutronOVSAgentLoggingBurstLimit**
  - **NeutronOVSAgentLoggingLocalOutputLogBase**
- Options for ML2/OVN back ends:
  - **NeutronOVNLoggingRateLimit**
  - **NeutronOVNLoggingBurstLimit**
  - **NeutronOVNLoggingLocalOutputLogBase** (BZ#1990357)

- This update fixes a bug that prevented the ceilometer-agent-compute service from collecting libvirt-related metrics.
  Previously, the libvirt service started after the ceilometer-agent-compute service, which resulted in "Permission denied" failures and loss of metrics data. Now the libvirt service starts before the ceilometer-agent-compute service and the service can properly collect metrics. (BZ#2092088)

- Previously, Red Hat Ceph Storage nodes were incorrectly configured to consume OpenStack high availability, advanced-virt, and fast-datapath repos during Leapp upgrades. The previous bug fix for this issue introduced an override that caused role-based parameters to work incorrectly.
  With this update, the role-based parameter implementation is fixed and the correct repositories are enabled for Red Hat Ceph Storage nodes. This update fixes the issue in Red Hat OpenStack Platform environments 16.2 and later that use the Red Hat Ceph Storage role. (BZ#2094377)

- This update fixes a bug that prevented the ceilometer-agent-compute service from polling for CPU metrics on Compute nodes. (BZ#2103970)

- Before this update, in overcloud deployments that enabled the Block Storage (cinder) backup service, a stack update affecting the Block Storage configuration did not restart the Block Storage service. This caused the Block Storage service to use the old configuration. With this update, the stack update procedure ensures that both the Block Storage backup service and the Block Storage service restart when the Block Storage configuration changes. This ensures that the Block Storage service always uses the latest configuration. (BZ#2106647)

- After you upgrade your host from Red Hat Enterprise Linux (RHEL) 7.9 to RHEL 8.4, you can specify additional packages to install in your environment by using the BaseTripeloPackages variable. With this feature, you can customize the base packages that your deployment requires on specific roles. For more information, see Customizing the base packages after a Leapp upgrade. (BZ#2123646)

- This update fixes a bug that prevented the Telemetry service (ceilometer) from polling the Object Storage service (swift) for metrics. The Telemetry service now polls the Object Storage service correctly. (BZ#2126786)

- This update fixes a bug that caused ceilometer-agent-ipmi to write log data inside the container namespace instead of on the host as expected.
  The improper placement of the content inside the container increased the container size, prevented proper log rotation, and resulted in loss of the log data when the container was deleted or rebuilt.

  Now ceilometer-agent-ipmi writes the logs to the host in **/var/log/containers/ceilometer/** as expected. (BZ#2138395)

- With this update, the collectd processes plugin has been removed from the default list of plugins. Loading the plugin can cause flooding issues and does not provide value when running in a containerized environment as it will only see the collectd and sensubility processes rather than the expected system processes. (BZ#2101944)

Changes to the openstack-tripleo-puppet-elements component:

- Before this update, some deployment or scale-up operations with affected builds failed due to a missing **tuned-profiles-cpu-partitioning** package on nodes that used the **overcloud-minimal** baremetal image, for example, Red Hat Ceph Storage nodes. With this update, the **tuned-profiles-cpu-partitioning** package is included in the **overcloud-minimal** image. This update restores the deployment and scale-up functionality for nodes that use the **overcloud-minimal** image. (BZ#2138203)

Changes to the openstack-tripleo-validations component:

- Before this update, the oslo-config-validation falsely reported errors with the "key_manager" and "barbican" sections of Block Storage (cinder) configuration. With this update, oslo-config-validation no longer falsely reports Block Storage configuration errors. (BZ#2112918)

Changes to the python-networking-ovn component:

- This update fixes a bug that causes connectivity loss after certain updates to RHOSP 16.2.2 and 16.2.3. If you are planning to update to a RHOSP 16.2 release, update to RHOSP 16.2.4 to avoid connectivity loss.
  The bug is triggered by a database schema change in OVN 21.12, which is introduced in RHOSP 16.2.2. and 16.2.3. OVN 21.12 contains a new column that is not present in earlier versions. OVN database schema changes should not cause a problem in OpenStack, but this particular change is affected by a bug.

In particular, instance connectivity is lost for a variable amount of time (from 20 seconds to 3 minutes) when you run the following command:

```
$ openstack overcloud external-update run --stack overcloud --tags ovn
```

To avoid the bug, do not update to RHOSP 16.2.2. or 16.2.3. Update to RHOSP 16.2.4 instead. (BZ#2111871)

Changes to the tripleo-ansible component:

- You can now specify a different Controller group name than the default, **Controller,** by using the `controller_group_name ` variable. (BZ#2037996)

4.3. RHEA-2022:4793 – RELEASE OF COMPONENTS FOR OSP 16.2.3

Changes to the collectd component:

- Libvirt was updated and as a result provides more metrics. Due to API changes, collectd was incompatible, resulting in a potential application crash. With this update, collectd was adjusted to provide hugepage usage via the virt plugin. The hugepages metric is exposed via the virt plugin and collectd was adjusted to match the API changes so that the application no longer crashes when pulling the virt metrics. (BZ#2038881)

Changes to the openstack-ironic component:

- Before this update, the RHOSP Bare Metal service (ironic) could lose its connection to the remote Redfish baseboard management controller (BMC) resulting in the bare metal node entering a maintenance state and with its power status changing to **None**. Depending on a site's environmental factors, some or all of the bare metal nodes could be in this unwanted maintenance state for an extended period of time.
  Transient network connectivity issues caused by high packet loss to the BMC caused connection caching issues when using Redfish. In cases where a session token needed to be renegotiated, the cached session object was never invalidated and connectivity was lost to the BMC.

  With this update, the Bare Metal service now initializes an entirely new cached session with a remote Redfish BMC when connectivity or authentication issues are detected. Additionally, this allows updated credentials to be leveraged should the bare metal node's BMC passwords be changed after initial configuration. (BZ#2064767)

Changes to the openstack-neutron component:

- Open vSwitch (OVS) does not support offloading OpenFlow rules that have the **skb_priority**, **skb_mark**, or output queue fields set. Those fields are needed to provide quality-of-service (QoS) support for virtio ports.
  If you set a minimum bandwidth rule for a virtio port, the Neutron Open vSwitch agent marks the traffic of this port with a Packet Mark Field. As a result, this traffic cannot be offloaded, and it affects the traffic in other ports. If you set a bandwidth limit rule, all traffic is marked with the default 0 queue, which means no traffic can be offloaded.

  As a workaround, if your environment includes OVS hardware offload ports, disable the packet marking in the nodes that require hardware offloading. After you disable the packet marking, it will not be possible to set rate limiting rules for virtio ports. However, differentiated services code point (DSCP) marking rules will still be available.

  In the configuration file, set the **disable_packet_marking** flag to **true**. After you edit the configuration file, you must restart the **neutron_ovs_agent** container. For example:

```
$ cat `/var/lib/config-data/puppet-
generated/neutron/etc/neutron/plugins/ml2/openvswitch_agent.ini`
  [ovs]
  disable_packet_marking=True
```

(BZ#2092946)

Changes to the openstack-tripleo-heat-templates component:

- In Red Hat OpenStack Platform (RHOSP) 16.2.3, the DM-Multipathing redundancy configuration for the Block Storage service (cinder) is now automated. (BZ#2047705)
- Before this update, during the replacement of a DCN node, the etcd service on the replacement node failed to start and caused the cinder-volume service on that node to fail. This failure was caused by the replacement for a DCN node attempting to start the etcd service as if it were bootstrapping a new etcd cluster, instead of joining the existing etcd cluster.
  With this update, a new parameter has been added, **EtcdInitialClusterState**. When **EtcdInitialClusterState** is set to **existing**, the DCN node starts etcd correctly, which causes the cinder-volume service to run successfully. (BZ#2055409)
- Before this update, if the **CinderDefaultVolumeType** parameter contained a non-alphanumeric character, the parameter validation failed. As a result, overcloud deployments and updates to existing deployments failed. With this update, **CinderDefaultVolumeType** validation supports non-alphanumeric characters. (BZ#2073096)
- In Red Hat OpenStack Platform (RHOSP) 16.2.3, there is RHOSP Orchestration service (heat) support for Dynamic Device Personalization (DDP) used in Intel Ethernet 800 Series NICs, such as the Columbiaville E810. For more information, see the Network Functions Virtualization Planning and Configuration Guide. (BZ#1967680)
- Before this update, attempts to stop the cinder-volume service running in active-active mode resulted in the **failed** state. The cause for these failures was that the service was not allowing sufficient time to properly stop itself.
  With this update, the time period allocated to stopping the service has been extended, and the service moves to the **inactive** state when you stop it. (BZ#2056918)
- This enhancement adds a new heat parameter, **FsAioMaxNumber**, that limits the number of parallel AIO requests to 1048576 by default. (BZ#2061697)

Changes to the openstack-tripleo-puppet-elements component:

- With RHOSP 16.2.3, fencing agents that depend on potentially insecure protocols, such as Telnet, have been removed from the default overcloud images. Overcloud images now include the following fencing agents by default:
- fence-agents-redfish
- fence-agents-ipmilan
- fence-agents-kdump
- fence-agents-rhevm
- fence-agents-compute
- sbd

- fence-agents-sbd (BZ#2081357)

Changes to the puppet-rsyslog component:

- Before this update, the **RsyslogElasticsearchSetting** parameter did not support the array format required to set multiple servers. With this update, the **RsyslogElasticsearchSetting** parameter now accepts values with multiple items. You can define multiple Elasticsearch servers by using the following array format: **server: ["elasticsearch1", "elasticsearch2"]** (BZ#2015699)

4.4. RHSA-2022:1001 – RELEASE OF COMPONENTS FOR OSP 16.2.2

Changes to the collectd-libpod-stats component:

- In cases where high CPU use was monitored in a multi-core system, the calculation for CPU use was inaccurate.
  With this update, the calculation of CPU use in a multi-core scenario is now accurate. The latest STF dashboards have been adjusted to incorporate this update. (BZ#2006970)

Changes to the openstack-cinder component:

- This update fixes a bug that caused some virtual machine (VM) bootup operations to fail when multiple VMs were booted simultaneously from an image.
  Previously, the Block Storage service (cinder) GPFS SpectrumScale driver did not correctly detect when the storage back end supported copy-on-write (COW) mode. The driver disabled COW features such as the ability to rapidly create volumes from an image. When booting multiple instances simultaneously from an image, some instances timed out when copying the image to its boot volume.

  The GPFS Block Storage service driver now properly detects when the storage back end supports COW mode. You can now spawn multiple VM instances simultaneously. (BZ#2022018)

- Before this update, when creating a snapshot with PowerMaxOS 5978.711, REST experienced a payload response change and caused the device label to modify its format. The underlying data from the solutions enabler changed and no longer contained a colon character (:). This caused an IndexError exception in the PowerMax Driver:

  ```
  IndexError: list index out of range
  ```

  With this update, the problem is resolved in PowerMaxOS 5978.711 and later. (BZ#2003762)

- This update fixes a bug that omitted details from the output of the **openstack volume backup list** command when the output exceeded 1000 lines. (BZ#2006556)

- Before this update, the NFS driver blocked attempts to delete OpenStack Storage snapshots in the error state, which prevented users from removing broken snapshot DB entries. With this update, the restriction is removed so that you can clean up failed snapshots. (BZ#2022121)

- This release includes an update of the Block Storage service (cinder) driver for HPE storage arrays. With the updated driver you can use the iSCSI protocol with HPE's Primera products for Primera version 4.2 and later. (BZ#2029943)

Changes to the openstack-neutron component:

- With this update, you can set QoS maximum bandwidth limit, egress direction rules on hardware-offloaded ports in a ML2/OVS deployment. To set the policy, use the normal QoS policy/rules methods.
  The back end uses **ip link** commands to enforce the policy instead of the normal OVS QoS engine, because the OVS **meter** action cannot be offloaded. See meter action is not offloaded. (BZ#1971545)

Changes to the openstack-tripleo-common component:

- Before this update, the dnsmasq healthcheck failed even when dnsmasq ran correctly. The healthcheck failed because it used the dnsmasq user rather than the root user, and did not have access to the /**proc** files. This resulted in incorrect systemd journal messages and failures when validations were enabled. With this update, the dnsmasq healthcheck is disabled because it is of limited use and it is being phased out in later releases. The dnsmasq container is now marked as healthy as long as it is running. (BZ#1961237)

Changes to the puppet-ironic component:

- With RHOSP 16.2.2, you can configure your undercloud to support both PXE and iPXE boot modes when your architecture includes both POWER (ppc64le) and x86_64 UEFI nodes. For more information, see Configuring a multiple CPU architecture overcloud . (BZ#1982489)

Changes to the puppet-tripleo component:

- This update corrects an error that prevented the proper use of the Block Storage service (cinder) **powermax_port_groups** parameter. (BZ#2027759)

Changes to the python-dracclient component:

- This update fixes a bug that caused unintended results when converting Dell EMC PERC H755 RAID controller physical disks to non-RAID mode.
  The conversion erroneously created RAID-0 virtual disks and moved them to the **Online** state, consuming a physical disk.

  The RAID-0 virtual disks are no longer created during the conversions. (BZ#2010246)

Changes to the python-os-brick component:

- This update fixes a bug that prevented connections to the Ceph storage backend with Ceph client release 15.2.0 (Octopus) and later, affecting Red Hat Ceph Storage 5.0 and later.
  A temporary configuration file generated to enable a Ceph connection did not include a '[global]' section marker. This update adds the '[global]' section marker to the temporary file.

  The section marker was introduced in Ceph client release 0.94.0 (Hammer). Starting with the Octopus release, Ceph requires the presence of the marker. This fix is backward compatible to Red Hat Ceph Storage 4.x. (BZ#2024684)

Changes to the python-tripleoclient component:

- Before this update, you could only configure the **net_config_override** parameter in **undercloud.conf** to point to a file in JSON format. With this update, you can now configure the **net_config_override** parameter in **undercloud.conf** to point to a file in either YAML or JSON format. (BZ#1956785)

Changes to the tripleo-ansible component:

- With this technology preview update, you can set the following parameters to configure OVS PMD automatic load balance:
  OvsPmdAutoLb: Enable/disable the OVS DPDK PMD Auto Load Balance feature. Values: true or false. OVS DPDK uses the default value of false.

  OvsPmdLoadThreshold: Set the minimum PMD thread load threshold for OVS DPDK PMD Auto Load Balance feature. Set a value from 0 to 100 to specify the minimum PMD thread load threshold (% of used cycles) of any non-isolated PMD threads when a PMD Auto Load Balance might be triggered.

  OvsPmdImprovementThreshold: Set PMD load variance improvement threshold for OVS DPDK PMD Auto Load Balance feature. Set a value from 0 to 100 to specify the minimum evaluated percentage improvement in load distribution across the non-isolated PMD threads that allows a PMD Auto Load Balance to occur.

  OvsPmdRebalInterval: Set PMD auto load balancing interval for OVS DPDK PMD Auto Load Balance feature. Set a value from 0 to 20,000 to specify the minimum time (in minutes) between 2 consecutive PMD Auto Load Balancing iterations. (BZ#1952060)

- With this update, you can use the **openstack undercloud backup** command with the **--db-only** option to create a backup of the database that runs on the undercloud node. You can use that backup to recover the state of the database in the event that it becomes corrupted during the upgrade process. (BZ#1977392)

4.5. RHSA-2022:0995 – RELEASE OF COMPONENTS FOR OSP 16.2.2

Security Fix(es):

- openstack-tripleo-heat-templates: data leak of internal URL through keystone_authtoken (CVE-2021-4180)
  For more details about the security issue(s), including the impact, a CVSS score, and other related information, refer to the CVE page(s) listed in the References section.

Bug Fix(es):

- Disable the advanced-virt-for-rhel-8 repository before you install Red Hat OpenStack Platform (RHOSP) 16.2, update from RHOSP 16.2 to a newer maintenance release, or upgrade from 16.1 to 16.2.
  RHOSP hosts do not require the advanced-virt-for-rhel-8 repository. If you do not disable it, dependency issues cause the installation, update, or upgrade to fail. The dependency failures happen because the advanced-virt-for-rhel-8-x86_64-rpms repository is based on RHEL 8.latest, which does not work with RHEL 8.4.

  As a workaround, disable the repositories. Perform the steps appropriate for your installation, update, or upgrade scenario.

- Scenario: new 16.2 installation or update from 16.2 to later version of 16.2.
  $ subscription-manager repos --disable advanced-virt-for-rhel-8-x86_64-rpms

  $ dnf module disable -y virt:av

  $ dnf module enable -y virt:rhel

- Scenario: upgrade from 16.1→16.2.
  $ subscription-manager repos --disable advanced-virt-for-rhel-8-x86_64-rpms

$ dnf module disable -y virt:8.2

$ dnf module enable -y virt:rhel (BZ#2027787)

- Red Hat OpenStack Platform (RHOSP) now supports the correct method of updating OVN. For more information, see Optional: Updating the ovn-controller container on all overcloud servers . (BZ#2050154)

- The Block Storage service (cinder) can now use an external NFS share to perform image format conversion of Image service (glance) images on the overcloud Controller nodes. Using this functionality prevents the space on the node from being completely filled during a conversion operation.
  See Configuring an external NFS share for conversion . (BZ#1886762)

- Before this update, the leapp process failed when the following statements were true:

- The system is deployed with TLS-everywhere (TLSe).

- The system uses the deprecated **authconfig** utility to configure authentication on your system.

- The command "leapp answer --section authselect_check.confirm=True" has not been run before triggering the leapp upgrade.
  With this update, the leapp process completes successfully without the need to run the command "leapp answer --section authselect_check.confirm=True" before triggering the leapp upgrade command. (BZ#1978228)

- You can now deploy the CephFS NFS gateway (ganesha) on the external network instead of the default dedicated StorageNFS network. (BZ#1999725)

- Before this update, missing parameters caused HAproxy logs to not be forwarded when rsyslog was configured using OpenStack Orchestration. With this update, all required parameters are present, which means that HAProxy logs are included in logs collected and forwarded. (BZ#2026290)

- This update adds an NTP validation step to deployments. To include the validation step in your deployment, include the argument **--ntp-server <ntp_server_name>** in the **openstack deploy** command. Replace <ntp_server_name> with the name of a valid, reachable NTP server.
  If director cannot reach the specified NTP server, the deployment fails. This validation prevents synchronization failures from occurring later in the deployment. (BZ#2034189)

- This update fixes a bug that caused the **metrics_qdr** service to appear in the HAProxy configuration of an Service Telemetry Framework (STF) deployment on a distributed compute node (DCN) node. (BZ#2038897)

- In Red Hat OpenStack Platform (RHOSP) deployments that use the Modular Layer 2 plug-in with the Open vSwitch (ML2/OVS) mechanism driver, there is currently a known issue where the Orchestration service (heat) parameter, **NeutronL3AgentAvailabilityZone** does not set the relevant Neutron L3 agent parameter correctly.
  Workaround: use a custom hieradata statement to set this value. In the example that follows, replace **[ROLE]** with the composable role name that is appropriate for your site.

  Example

```
[ROLE]ExtraConfig:
  neutron::agents::l3::availability_zone: role_availability_zone
```

For more information, see Puppet: Customizing hieradata for roles in the *Advanced Overcloud Customization* guide. (BZ#1983748)

- This update adds the **CollectdContainerAdditionalCapAdd** variable to the deployment tool. This variable is a comma separated list of additional collectd container capabilities. You can use it to add capabilities to the collectd container. (BZ#1984555)

- Sometimes a leapp actor blocks the leapp process during an upgrade. You can now remove actors with the new, role-specific **LeappActorsToRemove** parameter. (BZ#1984875)

- This enhancement updates the default value of the parameter **NovaLiveMigrationPermitAutoConverge** to enabled. When the parameter is enabled, the instance CPU is slowed down until the memory copy process is faster than the instance's memory writes when the migration performance is slow and might not complete. (BZ#1999324)

- This enhancement enables the experimental **rsyslog reopenOnTruncate** setting to ensure that rsyslog immediately recognizes when a logrotation happens on a file. The setting affects every service configured to work with rsyslog.
  When **rsyslog reopenOnTruncate** is disabled, rsyslog waits for a log file to fill to its original capacity before consuming any additional logs. (BZ#1949675)

- Before this update, the Ceph Dashboard could be exposed only on the provisioning network or on a dedicated composable network. Now, Red Hat OpenStack Platform director processes the **CephDashboardNetwork** parameter to define the frontend network used by operators, and limits access to the Ceph Dashboard and monitoring stack to this network. (BZ#1969411)

4.6. RHEA-2021:3483 – RELEASE OF COMPONENTS FOR OSP 16.2

Changes to the distribution component:

- The **python-networking-fujitsu** package is not included with RHOSP 16.2. (BZ#1906028)

Changes to the openstack-cinder component:

- Typically, when you create an encrypted volume from a snapshot of an encrypted volume, the source volume is the same size or smaller than the destination volume.
  In previous releases, if you created an encrypted volume from a snapshot of an encrypted volume, and the destination volume was close to or equal to the size of the source volume, the Block Storage service (cinder) silently truncated the data in the new destination volume.

  With this release, the Block Storage service calculates the size of the destination volume to include the current size of the encryption header, which eliminates the data truncation. (BZ#1772531)

- Before this update, the NetApp SolidFire driver would create a duplicate volume when the API response is lost due to a connection error and the driver retries the API request. This occurred when the SolidFire back end successfully received and processed a create volume operation, but failed to deliver the response back to the driver. This update resolves the issue by:

  1. Checking if the volume name already exists in the backend before trying to create it. If a volume is found, an exception is raised and the process is aborted.

  2. Checking for volume creation right after a read timeout is detected, to prevent invalid API calls.

3. Adding the ´sf_volume_create_timeout´ option to the SolidFire driver, to allow users to set the appropriate timeout value for their environment. (BZ#1939394)

- The Unisphere for PowerMax REST endpoints have changed from 91 to 92. This update changes how URIs are created, to allow for full coverage of all possible Unisphere REST API endpoints. (BZ#1913671)

- This update adds Challenge Handshake Authentication Protocol (CHAP) support to the Dell EMC PowerStore driver. PowerStore can now be used with enabled CHAP as a storage back end. (BZ#1935154)

- Before this update, if PowerStore ports were configured for multiple purposes, such as for iSCSI or Replication, the driver reported that it could find no accessible iSCSI targets. This was because the REST filter was wrong. This update fixes the PowerStore iSCSI targets filtering. (BZ#1953749)

- Before this update, when iSCSI or FC targets, such as ESXi, were not connected to the RHOSP host, the attach volume operation waited until it timed out. With this update, a new option to support ports filtering has been added to the Dell EMC XtremIO driver for the Block Storage service (cinder). (BZ#1956370)

Changes to the openstack-glance component:

- Before this update, RBD performance was degraded when multiple instances were launched simultaneously. This was due to the Image service starting multiple threads to perform the same copying operation. This update resolves the issue. (BZ#1851051)

- This update fixes an Image service (glance) configuration error that prevented users from creating a virtual machine with watchdog by setting flavor metadata. (BZ#1851797)

Changes to the openstack-ironic-python-agent component:

- In this release, EFI bootloader assets for whole-disk images are preserved during deployment, so the shim bootloader is no longer overwritten. This ensures that Secure Boot is switched on after deployment. (BZ#1961784)

Changes to the openstack-manila component:

- Previously, the Shared File Systems service (manila) API that brings external shares into service management did not check for duplicated export locations. An existing share brought into the service multiple times results in an inconsistent state.
  With this release, the API evaluates the export locations of known or existing shares before allowing external shares to be managed, and prevents existing shares from being erroneously brought into the Shared File Systems service again. (BZ#1849843)

- When multiple storage backends are configured on the Shared File Systems service (manila), each storage back end might support a different storage protocol. Before this update, the Shared File Systems service scheduler did not consider the storage protocol and capability of the shared storage back ends when deciding where to place them, which caused share provisioning to fail. With this update, the Shared File Systems service scheduler now automatically considers the share type extra specs with the storage protocol, which makes it possible use the **CapabilitiesFilter** to compare storage back end capabilities and provision shares successfully. (BZ#1888105)

Changes to the openstack-manila-ui component:

- Before this update, the Shared File Systems service (manila) dashboard had dynamic form elements whose names could potentially cause the forms to become unresponsive. This meant that the creation of share groups, share networks, and shares within share networks did not work. With this update, dynamic elements whose names could be problematic are encoded, which means that creation of share groups, share networks, and shares within share networks functions normally. (BZ#1974979)

Changes to the openstack-nova component:

- Before this update, when you resized or migrated an instance that had a vGPU flavor you needed to rebuild the instance manually to re-allocate the vGPU resources. With this update, instances with a vGPU flavor are automatically re-allocated the vGPU resources after resize and cold migration operations. (BZ#1844372)

- When an instance is created, the Compute service (nova) sanitizes the instance display name to generate a valid hostname when DNS integration is enabled in the Networking service (neutron).
  Before this update, the sanitization did not replace periods ('.') in instance names, for example, 'rhel-8.4'. This could result in display names being recognized as Fully Qualified Domain Names (FQDNs) which produced invalid hostnames. When instance names contained periods and DNS integration was enabled in the Networking service, the Networking service would reject the invalid hostname resulting in a failure to create the instance and a HTTP 500 server error from the Compute service.

  With this update, periods are now replaced by hyphens in instance names to prevent hostnames being parsed as FQDNs. You can continue to use free-form strings for instance display names. (BZ#1919855)

Changes to the openstack-tripleo-heat-templates component:

- This enhancement adds support for Entrust nShield HSM deployment in high availability mode with OpenStack Key Manager (barbican). (BZ#1714772)

- Images with null bytes take up a lot of space. With this release, you can enable sparse image upload and save space when you upload images. Sparse image upload is supported only with Ceph RBD. (BZ#1866741)

- In prior releases, you could not delete resources with dependencies, such as snapshots. With this release, you can delete resources with dependencies. (BZ#1884322)

- For Distributed Compute Node deployments which use storage, **dcn-hci.yaml** has been renamed to **dcn-storage.yaml** because DCN sites with storage have the option of not using HCI (Hyper-Converged Infrastructure). **dcn-hci.yaml** is deprecated but will remain in the environments directory for backwards compatibility. **dcn-hci.yaml** will be removed in Red Hat OpenStack platform 17. **dcn-storage.yaml** should be used in place of **dcn-hci.yaml**. (BZ#1868673)

- This enhancement improves the efficiency, performance, and execution time of deployment and update tasks for environments with a large number of roles. The logging output of the deployment process has been improved to include task IDs for better tracking of specific task executions, which can occur at different times. The task IDs can now be used to correlate timing and execution when troubleshooting executions. (BZ#1897890)

- With this enhancement, you can improve the performance of live migrations by using the following new parameters:

  - **NovaLiveMigrationPermitPostCopy** - When enabled, the instance is activated on the

destination node before migration is complete, and an upper bound is set on the memory that needs to be transferred, which improves the live migration of larger instances. This parameter is enabled by default.

- **NovaLiveMigrationPermitAutoConverge** - When enabled, if an on-going live migration is progressing slowly the instance CPU is throttled until the memory copy process is faster than the instance's memory writes. This parameter is disabled by default. To enable **NovaLiveMigrationPermitAutoConverge**, add the following configuration to an environment file:

```
parameter_defaults:
  ComputeParameters:
    NovaLiveMigrationPermitAutoConverge: true
```

(BZ#1920229)

Changes to the openstack-tripleo-validations component:

- This enhancement adds new validation for **tripleo-latest-packages-version**. This validation checks if the listed **tripleo** packages are up to date with repositories. (BZ#1926725)

Changes to the puppet-collectd component:

- Before this update, the **PluginInstanceFormat** parameter for **collectd** could specify only one of the following values: 'none', 'name', 'uuid', or 'metadata'. After this update, the **PluginInstanceFormat** parameter for **collectd** can now specify more than one value, which results in more information being sent in the **plugin_instance** label of **collectd** metrics. (BZ#1938568)

Changes to the python-glance-store component:

- Before this update, writing an image to RBD could be very slow. This update improves the process for writing an image to RBD, which improves the time it takes for images to be written to RBD. (BZ#1690726)

Changes to the python-networking-ovn component:

- This update fixes an issue that caused Networking service (neutron) agents, such as Networking service DHCP, to fail when they tried to create resources in OVN because ML2/OVN prevented RPC workers from connecting the OVN southbound database. (BZ#1972774)

- This update fixes an issue that caused Networking service (neutron) agents, such as Networking service DHCP, to fail when they tried to create resources in OVN. This was caused by residual data left in the OVN databases when QoS rules were created for floating IPs.
  This update eliminates the residual data and fixes the problem. (BZ#1978158)

Changes to the python-os-brick component:

- Before this update, some exceptions were not being caught during connections to iSCSI portals, such as failures in **iscsiadm -m session**. This caused **_connect_vol** threads to abort unexpectedly in some failure patterns, which caused subsequent steps to hang while waiting for results from **_connect_vol** threads. This update ensures that any exceptions during connections to iSCSI portals are handled correctly in the **_connect_vol** method, to avoid unhandled exceptions during connecting to iSCSI portals, and unexpected aborts that have no updated thread results. (BZ#1923975)

Changes to the python-oslo-config component:

- This enhancement adds the type **HostDomain**. **HostDomain** is the same as **HostAddress** with the added support of the underscore character - RFC1033. Systems such as DomainKeys and service records use the underscore. The Compute service can use the **HostDomain** type to define **live_migration_inbound_addr**. (BZ#1868940)

Changes to the python-tripleoclient component:

- Before this update, the validation variable in one code path was referenced but never assigned, which resulted in an unhandled exception during validation. This has been fixed. (BZ#1959853)

Changes to the tripleo-ansible component:

- Starting with Red Hat Enterprise Linux (RHEL) version 8.3, support for the Intel Transactional Synchronization Extensions (TSX) feature is disabled by default. Currently, this causes instance live migration to fail when migrating from hosts where the TSX kernel argument is enabled to hosts where the TSX kernel argument is disabled.
  This impact applies only to Intel hosts that support the TSX feature. For more information about the CPUs that are affected by this issue, see Affected Configurations.

  For more information, review the following Red Hat Knowledgebase solution Guidance on Intel TSX impact on OpenStack guests. (BZ#1975240)

- In Red Hat OpenStack Platform 16.2, a technology preview is available that supports Precision Time Protocol (PTP) with Timemaster. (BZ#1825895)

- Before this fix, grub2 tooling wrote kernel argument changes to /boot/grub2/grubenv. This file was not available to UEFI boot systems, and caused kernel argument changes not to persist across reboots on UEFI boot nodes.
  This fix changes both the /boot/grub2/grubenv file and the /boot/efi/EFI/redhat/grubenv files when you make kernel argument changes. As a result, RHOSP director now applies persistent Kernel argument changes for UEFI boot nodes. (BZ#1987092)

- During stack update the **KernelArgs** could be modified or appended. A reboot of the affected nodes needs to be performed manually.
  For example, if the current deployment has the following configuration, it is possible to change **hugepages=64**, or add or remove arguments during the stack update:

  ```
  `KernelArgs: "default_hugepagesz=1GB hugepagesz=1G hugepages=32 intel_iommu=on iommu=pt isolcpus=1-11,13-23"
  ```

  For example:

  ```
  KernelArgs: "default_hugepagesz=1GB hugepagesz=1G hugepages=64 intel_iommu=on iommu=pt isolcpus=1-24"
  KernelArgs: "isolcpus=1-11,13-23"
  ```

  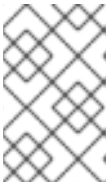

  NOTE

  Complete removal of **KernelArgs** during the update is not supported. Also **KernelArgs** could be newly added as well to an existing overcloud node, however the reboot would be triggered in this case. (BZ#1900723)

Changes to the validations-common component:

- Before this update, execution of the validation package **check-latest-packages-version** was slow. This update resolves the issue. (BZ#1942531)

- This enhancement improves the performance and application of the **check-latest-packages-version** validation. (BZ#1926721)

- Before this update, validation results were not being logged and validation artifacts were not being collected as the permissions required to access the requested logging directory were not granted. This update resolves the issue, and validation results are successfully logged and validation artifacts are collected. (BZ#1910508)

- Before this update, Ansible redirected output to all registered non-stdout callback plug-ins by default, which resulted in VF callbacks processing information from other processes using **ansible runtime**. This issue has been resolved and the output of other processes is no longer stored in the validations logging directory. (BZ#1960185)