Red Hat OpenStack Platform 17.1

Release notes

Release details for Red Hat OpenStack Platform 17.1

Last Updated: 2024-09-06

Red Hat OpenStack Platform 17.1 Release notes

Release details for Red Hat OpenStack Platform 17.1

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 for Red Hat OpenStack Services on OpenShift (RHOSO) or earlier releases of Red Hat OpenStack Platform (RHOSP). When you create an issue for RHOSO or RHOSP documents, the issue is recorded in the RHOSO Jira project, where you can track the progress of your feedback.

To complete the Create Issue form, ensure that you are logged in to Jira. If you do not have a Red Hat Jira account, you can create an account at https://issues.redhat.com.

1. Click the following link to open a **Create Issue** page: Create Issue
2. 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.
3. Click **Create**.

CHAPTER 1. INTRODUCTION

1.1. ABOUT THIS RELEASE

This release of Red Hat OpenStack Platform (RHOSP) is based on the OpenStack "Wallaby" 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 "Wallaby" release itself are available at the following location: https://releases.openstack.org/wallaby/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 9.2 Extended Update Support (EUS).

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. 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 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.

1.6. PRODUCT CERTIFICATION CATALOG

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

1.7. COMPUTE DRIVERS

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

This release of 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.

1.8. CONTENT DELIVERY NETWORK (CDN) REPOSITORIES

This section describes the repositories required to deploy Red Hat OpenStack Platform (RHOSP) 17.1.

You can install RHOSP 17.1 through the Content Delivery Network (CDN) by using **subscription-manager**.

For more information, see Planning your 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.

1.8.1. Undercloud repositories

You run Red Hat OpenStack Platform (RHOSP) 17.1 on Red Hat Enterprise Linux (RHEL) 9.2.

NOTE

If you synchronize repositories with Red Hat Satellite, you can enable specific versions of the Red Hat Enterprise Linux repositories. However, the repository remains the same despite the version you choose. For example, you can enable the 9.2 version of the BaseOS repository, but the repository name is still **rhel-9-for-x86_64-baseos-eus-rpms** despite the specific version you choose.

WARNING

Any repositories except the ones specified here are not supported. Unless recommended, do not enable any other products or repositories except 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 9 for x86_64 - BaseOS (RPMs) Extended Update Support (EUS)	**rhel-9-for-x86_64-baseos-eus-rpms**	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 9 for x86_64 - AppStream (RPMs)	**rhel-9-for-x86_64-appstream-eus-rpms**	Contains Red Hat OpenStack Platform dependencies.
Red Hat Enterprise Linux 9 for x86_64 - High Availability (RPMs) Extended Update Support (EUS)	**rhel-9-for-x86_64-highavailability-eus-rpms**	High availability tools for Red Hat Enterprise Linux. Used for Controller node high availability.
Red Hat OpenStack Platform for RHEL 9 (RPMs)	**openstack-17.1-for-rhel-9-x86_64-rpms**	Core Red Hat OpenStack Platform repository, which contains packages for Red Hat OpenStack Platform director.
Red Hat Fast Datapath for RHEL 9 (RPMS)	**fast-datapath-for-rhel-9-x86_64-rpms**	Provides Open vSwitch (OVS) packages for OpenStack Platform.

1.8.2. Overcloud repositories

You run Red Hat OpenStack Platform (RHOSP) 17.1 on Red Hat Enterprise Linux (RHEL) 9.2.

NOTE

If you synchronize repositories with Red Hat Satellite, you can enable specific versions of the Red Hat Enterprise Linux repositories. However, the repository remains the same despite the version you choose. For example, you can enable the 9.2 version of the BaseOS repository, but the repository name is still **rhel-9-for-x86_64-baseos-eus-rpms** despite the specific version you choose.

WARNING

Any repositories except the ones specified here are not supported. Unless recommended, do not enable any other products or repositories except 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 9 for x86_64 - BaseOS (RPMs) Extended Update Support (EUS)	**rhel-9-for-x86_64-baseos-eus-rpms**	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 9 for x86_64 - AppStream (RPMs)	**rhel-9-for-x86_64-appstream-eus-rpms**	Contains Red Hat OpenStack Platform dependencies.
Red Hat Enterprise Linux 9 for x86_64 - High Availability (RPMs) Extended Update Support (EUS)	**rhel-9-for-x86_64-highavailability-eus-rpms**	High availability tools for Red Hat Enterprise Linux.
Red Hat OpenStack Platform for RHEL 9 x86_64 (RPMs)	**openstack-17.1-for-rhel-9-x86_64-rpms**	Core Red Hat OpenStack Platform repository.
Red Hat Fast Datapath for RHEL 9 (RPMS)	**fast-datapath-for-rhel-9-x86_64-rpms**	Provides Open vSwitch (OVS) packages for OpenStack Platform.
Red Hat Ceph Storage Tools 6 for RHEL 9 x86_64 (RPMs)	**rhceph-6-tools-for-rhel-9-x86_64-rpms**	Tools for Red Hat Ceph Storage 6 for Red Hat Enterprise Linux 9.

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 9 for x86_64 - BaseOS (RPMs) Extended Update Support (EUS)	**rhel-9-for-x86_64-baseos-eus-rpms**	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 9 for x86_64 - AppStream (RPMs)	**rhel-9-for-x86_64-appstream-eus-rpms**	Contains Red Hat OpenStack Platform dependencies.
Red Hat Enterprise Linux 9 for x86_64 - High Availability (RPMs) Extended Update Support (EUS)	**rhel-9-for-x86_64-highavailability-eus-rpms**	High availability tools for Red Hat Enterprise Linux.
Red Hat OpenStack Platform for RHEL 9 x86_64 (RPMs)	**openstack-17.1-for-rhel-9-x86_64-rpms**	Core Red Hat OpenStack Platform repository.
Red Hat Fast Datapath for RHEL 9 (RPMS)	**fast-datapath-for-rhel-9-x86_64-rpms**	Provides Open vSwitch (OVS) packages for OpenStack Platform.
Red Hat Ceph Storage Tools 6 for RHEL 9 x86_64 (RPMs)	**rhceph-6-tools-for-rhel-9-x86_64-rpms**	Tools for Red Hat Ceph Storage 6 for Red Hat Enterprise Linux 9.

Ceph Storage node repositories

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

Name	Repository	Description of requirement
Red Hat Enterprise Linux 9 for x86_64 - BaseOS (RPMs)	**rhel-9-for-x86_64-baseos-rpms**	Base operating system repository for x86_64 systems.
Red Hat Enterprise Linux 9 for x86_64 - AppStream (RPMs)	**rhel-9-for-x86_64-appstream-rpms**	Contains Red Hat OpenStack Platform dependencies.
Red Hat OpenStack Platform Deployment Tools for RHEL 9 x86_64 (RPMs)	**openstack-17.1-deployment-tools-for-rhel-9-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-17.1-for-rhel-9-x86_64-rpms** repository.

Name	Repository	Description of requirement
Red Hat OpenStack Platform for RHEL 9 x86_64 (RPMs)	**openstack-17.1-for-rhel-9-x86_64-rpms**	Packages to help director configure Ceph Storage nodes. This repository is included with combined Red Hat OpenStack Platform and Red Hat Ceph Storage subscriptions. If you use a standalone Red Hat Ceph Storage subscription, use the **openstack-17.1-deployment-tools-for-rhel-9-x86_64-rpms** repository.
Red Hat Ceph Storage Tools 6 for RHEL 9 x86_64 (RPMs)	**rhceph-6-tools-for-rhel-9-x86_64-rpms**	Provides tools for nodes to communicate with the Ceph Storage cluster.
Red Hat Fast Datapath for RHEL 9 (RPMS)	**fast-datapath-for-rhel-9-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.

1.9. PRODUCT SUPPORT

The resources available for product support include the following:

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. You can access the following facilities through the Customer Portal:

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

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

Mailing Lists

You can join the **rhsa-announce** public mailing list to receive notification of security fixes for RHOSP and other Red Hat products.
Subscribe at https://www.redhat.com/mailman/listinfo/rhsa-announce.

1.10. UNSUPPORTED FEATURES

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

- Custom policies, which includes modification of **policy.json** or **policy.yaml** 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.
- Persistent memory for instances (vPMEM). You can create persistent memory namespaces only on Compute nodes that have NVDIMM hardware. Red Hat has removed support for persistent memory from RHOSP 17+ 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
- Virtualized control planes.

If you require support for any of these features, contact the Red Hat Customer Experience and Engagement team to discuss a support exception, if applicable, or other options.

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. BACKUP AND RESTORE

This section outlines the top new features related to backing up and restoring the Red Hat OpenStack Platform (RHOSP) undercloud and control plane nodes.

Snapshot and revert

The RHOSP snapshot and revert feature is based on the Logical Volume Manager (LVM) snapshot functionality and reverts an unsuccessful upgrade or update. Snapshots preserve the original disk state of your RHOSP cluster before performing an upgrade or an update. You can then remove or revert the snapshots depending on the results. If an upgrade completes successfully and you do not need the snapshots anymore, remove them from your nodes. If an upgrade fails, you can revert the snapshots, assess any errors, and start the upgrade procedure again. A revert leaves the disks of all the nodes exactly as they were when the snapshot was taken.

2.2. BARE METAL PROVISIONING

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

LVM thin provisioning

In RHOSP 17.1, the LVM volumes installed by the **overcloud-hardened-uefi-full.qcow2** whole disk overcloud image are now backed by a thin pool. By default, the volumes expand to consume the available physical storage, but they are not over-provisioned.

2.3. COMPUTE

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

Moving to Q35 default machine type

The default machine type for each host architecture is Q35 for new RHOSP 17 deployments. The Q35 machine type provides several benefits and improvements, including live migration of instances between different RHEL 9.x minor releases, and native PCIe hotplug, which is faster than the ACPI hotplug used by the i440FX machine type. You can still use the i440FX machine type.

Emulated virtual Trusted Platform Module (vTPM) devices for instances

You can use TPM to enhance computer security and provide a chain of trust for virtualization. The emulated vTPM is a software-based representation of a physical TPM chip. An administrator can provide cloud users the ability to create instances that have vTPM devices.

UEFI Secure Boot

Cloud users can launch instances that are protected with UEFI Secure Boot when the overcloud contains UEFI Secure Boot Compute nodes. For information about creating an image for UEFI Secure Boot, see Creating an image for UEFI Secure Boot . For information about creating a flavor for UEFI Secure Boot, see "UEFI Secure Boot" in Flavor metadata.

Ability to create instances that have a mix of dedicated and shared CPUs

You can now create flavors that have a mixed CPU policy to enable your cloud users to create instances that have a mix of dedicated (pinned) and shared (unpinned) CPUs.

VirtIO data path acceleration (VDPA) support for enterprise workloads

On RHOSP deployments that are configured for OVS hardware offload and to use ML2/OVN, and that have Compute nodes with VDPA devices and drivers and Mellanox NICs, you can enable your cloud users to create instances that use VirtIO data path acceleration (VDPA) ports. For more information, see Configuring VDPA Compute nodes to enable instances that use VDPA ports and Creating an instance with a VDPA interface .

Scheduler support for routed networks

On RHOSP deployments that use a routed provider network, you can now configure the Compute scheduler to filter Compute nodes that have affinity with routed network segments, and verify the network in placement before scheduling an instance on a Compute node. You can enable this feature by using the **NovaSchedulerQueryPlacementForRoutedNetworkAggregates** parameter.

2.4. DISTRIBUTED COMPUTE NODES (DCN)

This section outlines the top new features for Distributed Compute Nodes (DCN).

Framework for upgrades for Distributed Compute Node Architecture

In RHOSP 17.1.3, Red Hat now supports upgrading edge deployed architectures from 16.2 to 17.1 using the framework for upgrades workflow.

2.5. NETWORKING

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

Revert back to the OVS mechanism driver after failed migration to OVN

Starting in RHOSP 17.1.3 you can revert a failed or interrupted migration if you first follow the proper backup steps and revert instructions. The reverted OVS environment might be altered from the original. For example, if you migrate to the OVN mechanism driver, then migrate an instance to another Compute node, and then revert the OVN migration, the instance will be on the original Compute node. Also, a revert operation interrupts connection to the data plane.

HTTP/2 listener support for TLS-terminated load balancers

RHOSP 17.1.2 introduces support for TLS-terminated HTTP/2 listeners. HTTP/2 listeners enable you to improve the user experience by loading web pages faster and by employing the Application-Layer Protocol Negotiation (ALPN) TLS extension when load balancers negotiate with clients.
For more information about HTTP/2 listener support, see Creating a TLS-terminated load balancer with an HTTP/2 listener in *Configuring load balancing as a service* .

Migration to OVN mechanism driver

Migrations from the OVS mechanism driver to the OVN mechanism driver were not supported in RHOSP 17.0 because upgrades to RHOSP 17.0 were not supported.
OVN migration is now supported in RHOSP 17.1 GA.

You have the choice of migrating from ML2/OVS in 16.2 or 17.1. In most cases, Red Hat recommends upgrading to RHOSP 17.1 before migrating to ML2/OVN, because of enhanced functionality and improved migration functions in RHOSP 17.1.

Stateless security groups

This RHOSP release introduces support of the OpenStack stateless security groups API with the ML2/OVN mechanism driver. Stateless security groups are not supported by RHOSP deployments with the ML2/OVS mechanism driver.

A stateless security group can provide performance benefits because it bypasses connection tracking in the underlying firewall, providing an option to offload conntrack-related OpenFlow rules to hardware.

For more information about stateless security groups, see Configuring security groups.

Security group logging

To monitor traffic flows and attempts into and out of an instance, you can create packet logs for security groups. 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.

Quality of Service (QoS) for egress on hardware offloaded ports

Starting with RHOSP 17.1, in ML2/OVN deployments, you can enable minimum bandwidth and bandwidth limit egress policies for hardware offloaded ports. You cannot enable ingress policies for hardware offloaded ports. For more information, see Configuring the Networking service for QoS policies.

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

Starting in RHOSP 17.1, OVS PMD moves from technology preview to full support. 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.
  For more information, see Configuring DPDK parameters for node provisioning.

2.6. NETWORK FUNCTIONS VIRTUALIZATION

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

OVS and OVN TC Flower offload with Conntrack

In RHOSP 17.1, connection tracking (conntrack) hardware offloading is supported for ML2/OVS and ML2/OVN with TC Flower. For the conntrack module to offload openflow flows to hardware, you must enable security groups and port security on **switchdev** ports. For more information, see Configuring OVS TC-flower hardware offload .

2.7. SECURITY

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

FIPS 140-3 compatibility is now fully supported

You can now enable FIPS 140-3 compatibility mode with RHOSP.

SRBAC is now fully supported

You can now enable secure role-based access control in RHOSP.

2.8. STORAGE

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

Upgrade Red Hat Ceph Storage 5 to 6

Upgrading your Red Hat Ceph Storage cluster from version 5 to version 6 is now supported as a step in upgrading your RHOSP to RHOSP 17.1.2.
Upgrading directly from Red Hat Ceph Storage version 4 to version 6 is not supported. If you are currently using Red Hat Ceph Storage version 4, you must upgrade to Red Hat Ceph Storage version 5 before upgrading to Red Hat Ceph Storage version 6.

For more information about this procedure see, Framework for upgrades (16.2 to 17.1) .

Red Hat Ceph Storage 6

If you deploy greenfield RHOSP 17.1 with Red Hat Ceph Storage (RHCS), RHOSP is deployed with RHCS 6.1. RHCS 6 is also supported as an external Red Hat Ceph Storage cluster.

Red Hat Ceph Storage 7

RHOSP 17.1.3 adds support for RHCS 7 as an external Red Hat Ceph Storage cluster.

Availability zones for file shares

In RHOSP 17.1, cloud administrators can configure availability zones for Shared File Systems service (manila) back ends.

Manage/unmanage file shares

In RHOSP 17.1, cloud administrators can bring shares that are created outside the Shared File Systems service (manila) under the management of the Shared File Systems service and remove shares from the Shared File Systems service without deleting them. The CephFS driver does not support this feature. You can use this manage/unmanage functionality when commissioning, decommissioning, or migrating storage systems, or to take shares offline temporarily for maintenance.

Block Storage supports NVMe over TCP back ends

In RHOSP 17.1, the Block Storage service (cinder) supports NVMe over TCP (NVMe/TCP) drivers, for Compute nodes that are running RHEL 9.

Active-active configuration for the Block Storage backup service

In RHOSP 17.1, the Block Storage (cinder) backup service is deployed using an active-active configuration. For more information, see Deploying your active-active Block Storage backup service .

Other Block Storage backup service improvements

In RHOSP 17.1, the Block Storage (cinder) backup service supports the S3 back end and the zstd data compression algorithm. For more information, see Backup repository back-end configuration and Block Storage backup service configuration .

New Dell PowerFlex and PowerStore drivers

The Shared File Systems service (manila) now includes back-end drivers to provision and manage NFS shares on Dell PowerFlex storage systems, and NFS and CIFS shares on Dell PowerStore storage systems. The use of these drivers is supported when the vendor publishes certification on the Ecosystem Catalog.

2.9. UPGRADES AND UPDATES

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

Pre-update and post-update validations

In RHOSP 17.1.1, pre-update and post-update validations are now supported. With this enhancement, you can verify the requirements and functionality of your undercloud before you begin a minor update of your environment. You can then verify the overcloud functionality after you perform a minor update. For more information, see Validating RHOSP before the undercloud update and Validating RHOSP after the overcloud update in *Performing a minor update of Red Hat OpenStack Platform*.

Multi-RHEL

In RHOSP 17.1, you can upgrade a portion of your Compute nodes to RHEL 9.2 while the rest of your Compute nodes remain on RHEL 8.4. This is referred to as a Multi-RHEL environment. For more information about the benefits and limitations of a Multi-RHEL environment, see Planning for a Compute node upgrade in *Framework for upgrades (16.2 to 17.1)* .

2.10. TECHNOLOGY PREVIEWS

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

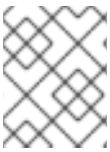

NOTE

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

Router flavors

The router flavors feature lets you define router flavors and use them to create custom virtual routers. For more information, see Creating custom virtual routers with router flavors.

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 17.1.3 MAINTENANCE RELEASE - MAY 22, 2024

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:

RHSA-2024:2727

Important: Red Hat OpenStack Platform 17.1 (python-gunicorn) security update

RHSA-2024:2729

Important: Red Hat OpenStack Platform 17.1 (etcd) security update

RHSA-2024:2730

Important: Red Hat OpenStack Platform 17.1 (collectd-sensubility) security update

RHSA-2024:2731

Moderate: Red Hat OpenStack Platform 17.1 (python-django) security update

RHSA-2024:2732

Moderate: Red Hat OpenStack Platform 17.1 (python-glance-store) security update

RHSA-2024:2733

Moderate: Red Hat OpenStack Platform 17.1 (openstack-ansible-core) security update

RHSA-2024:2734

Moderate: Red Hat OpenStack Platform 17.1 (python-urllib3) security update

RHSA-2024:2735

Moderate: Red Hat OpenStack Platform 17.1 (python-paramiko) security update

RHSA-2024:2736

Moderate: Red Hat OpenStack Platform 17.1 (openstack-tripleo-heat-templates and tripleo-ansible) security update

RHSA-2024:2737

Moderate: Red Hat OpenStack Platform 17.1 (python-openstackclient) security update

RHBA-2024:2738

Updated Red Hat OpenStack Platform 17.1 container images

RHBA-2024:2739

Updated Red Hat OpenStack Platform 17.1 container images

RHBA-2024:2740

Red Hat OpenStack Platform 17.1 RHEL 9 director images

RHBA-2024:2741

Red Hat OpenStack Platform 17.1 bug fix and enhancement advisory

RHBA-2024:2742

Red Hat OpenStack Platform 17.1 bug fix and enhancement advisory

RHSA-2024:2767

Important: Red Hat OpenStack Platform 17.1 (collectd-sensubility) security update

RHSA-2024:2768

Moderate: Red Hat OpenStack Platform 17.1 (python-paramiko) security update

RHSA-2024:2769

Moderate: Red Hat OpenStack Platform 17.1 (python-openstackclient) security update

RHSA-2024:2770

Moderate: Red Hat OpenStack Platform 17.1 (tripleo-ansible and openstack-tripleo-heat-templates) security update

3.1.2. Bug fixes

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

BZ#2222683

With this update, Red Hat support for Multi-RHEL is expanded to include the following deployment architectures:

- Edge (DCN)
- ShiftOnStack
- director Operator-based deployments

BZ#2229779

This update improved the OVN metadata agent process exception handling in an event when a child process is not able to start. With this update, the child process exception is logged and the main OVN metadata agent continues running.

BZ#2237866

Before this update, configuring caching parameters for ceilometer were not supported. With this update, for caching, ceilometer uses the **dogpile.cache.memcached** back end. If you manually disable caching, ceilometer uses the **oslo_cache.dict** back end.

BZ#2248873

With some versions of Pluggable Authentication Modules (PAM), the pam_loginuid module **/proc/self/loginuid** must be writable. This is not the case in the **sshd** container used for migrations. Migrations failed because SSH login between Compute hosts was failing. With this update, the **pam_loginuid** module has been removed from the PAM config and, as a result, SSH login between Compute hosts, and migrations, work again.

BZ#2249444

Before this update, libvirt-related containers on Compute nodes were needlessly restarted during deployment, update, and scaling operations, even when the container configurations were not changed.
With this update, libvirt-related containers that do not have configuration changes are not restarted during deployment, update, and scaling operations.

BZ#2249690

Before this update, after every successful Red Hat Ceph Storage adoption during upgrade, the **ceph-ansible** package was removed by default.
This update introduces the tag, **cleanup_cephansible**, to the task that removes **ceph-ansible**. You can use this tag with **--skip-tags** while running the adoption playbook to avoid removal.

BZ#2254036

Before this update, during a DCN FFU system upgrade of nodes on the setup with multiple stacks, the Red Hat Ceph Storage task **Set noout flag** might fail to run the ceph command on the right host.
After the update, a system upgrade on any node in a multi-stack setup now delegates the Red Hat Ceph Storage task **Set noout flag** to the relevant host, and the **ceph** commands are run on the specific cluster.

BZ#2254994

This update fixes a bug that caused unintended deletion of Load-balancing service (octavia) health monitor ports and OVN metadata ports.
Before this update, in RHOSP environments that used Load-balancing service health monitor ports from a previous version, running **neutron-db-sync-tool** sometimes randomly deleted those pre-existing ports or OVN metadata ports. This unintended port deletion resulted in loss of health monitor capacity or communication loss with the affected instances.

This update resolves the conflicts. It uses a new value, **ovn-lb-hm:distributed**, for the OVN Load-balancing service health monitor ports **device_owner** field. Old OVN Load-balancing service health monitor ports are automatically updated with this version.

BZ#2255324

Before this update, a director bug could disrupt or crash client workloads during updates or upgrades to any RHOSP 17.1 version. This bug affected deployments that enabled the RHOSP Shared File System service (manila) with the CephFS on NFS back ends.
With this update this issue has been addressed, and the Shared File System service (manila) operates properly when users set up "access rules" on their shares.

BZ#2257274

Before this update, when using jumbo frames for Networking service (neutron) tenant networks, a RHOSP Controller shutting down could sometimes cause the RHOSP Load-balancing service (octavia) management interface (**o-hm0**) to have its MTU reset to a small value, such as 1500 or 1450. This problem usually occurred when the RHOSP Controller was rebooted for the first time, or in a situation when the Controller was abruptly terminated. With this update, RHOSP director now ensures that Open vSwitch (OVS) is configured with the correct MTU when the **o-hm0** is created.

BZ#2259286

Before this update, the FFU procedure sometimes used the incorrect Red Hat Ceph Storage image when the documented upgrade procedure was not followed by the user.
With this update, the FFU procedure always uses the correct Red Hat Ceph Storage image. The **multi-rhel-container-image-prepare.py** script has been updated to use the correct defaults and version validation checks have been added to the FFU process.

BZ#2263552

This update fixes a bug that prevented load-balancing of traffic on some Load-balancing service (octavia) pool members on IPv6 networks in an ML2/OVN environment.
Before this update, if you added a second **listener+pool+member** to a pool, the pool entered an ERROR state and traffic within that pool was not load-balanced.

With this update, traffic is load-balanced to all members as expected.

BZ#2263916

This update safeguards against upgrades from RHOSP 16.2 to RHOSP 17.1 with libvirt configurations that could result in workload disruptions after the upgrade.
Before this update, if you performed an upgrade to RHOSP 17.1 from a RHOSP 16.2 environment that included a modular deployment of libvirt on Red Hat Enterprise Linux (RHEL) 8 or that ran libvirt UBI9 on RHEL 8, these configurations sometimes resulted in workload disruption.

With this update, the upgrade from RHOSP 16.2 to 17.1 fails if the RHOSP 16.2 environment includes a modular deployment of libvirt on Red Hat Enterprise Linux (RHEL) 8, or runs libvirt UBI9 on RHEL 8.

BZ#2266285

This update fixes a bug that prevented operation of the OVN Health Monitoring for Load-balancing service (octavia) on IPv6 networks in ML2/OVN deployments.
Previously, the OVN Health Monitors service did not correctly identify the **ONLINE** and **OFFLINE** statuses of back-end members.

With this update, load-balancing works as expected, and the OVN Health Monitors correctly identify the **ONLINE** and **OFFLINE** statuses of back-end members.

BZ#2278028

Before this update, during upgrades to RHOSP 17.1 and updates between minor versions of 17.1, Networking service (neutron) environments that used the ML2/OVN mechanism driver stopped receiving updates from the OVN databases for a period of time. When the Controller node that contained the RAFT leader was also updated, the mechanism driver would then receive the updates from the OVN databases.
With this update, this problem has been fixed. Now, during RHOSP updates and upgrades, the Networking service using the ML2/OVN mechanism driver operates properly.

3.1.3. Enhancements

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

BZ#1900663

With this update, Red Hat support for Framework for upgrades is expanded to include DCN deployments without storage at the edge.

BZ#1997638

With this update, Red Hat support for Framework for upgrades is expanded to include DCN deployments with storage at the edge.

BZ#2218000

With this update, you can now use the Bare Metal service (ironic) to boot an ISO image directly for use as a RAM disk. For more information, see Enabling ISO boot for bare-metal instances and Booting an ISO image directly for use as a RAM disk .

BZ#2224492

RHOSP 17.1.3 now supports a new aging mechanism for MAC addresses learned through the **localnet_learn_fdb** option that is included in Open Virtual Network (OVN) version 23.09. This new aging mechanism consists of two new options, **fdb_age_thhreshold** and **fdb_removal_limit**. The **fdb_age_thhreshold** option enables you to set the maximum time the learned MACs stay in the FDB table (in seconds).

The **fdb_removal_limit** option prevents OVN from removing a large number of FDB table entries all at one time.

When you use these new options with **localnet_learn_fdb**, you reduce the likelihood of performance and scalability issues caused by the FDB table from growing too big, a problem often experienced in RHOSP environments that have large provider networks.

BZ#2225163

A power save profile, cpu-partitioning-powersave, has been introduced in Red Hat Enterprise Linux 9 (RHEL 9), and is now available in Red Hat OpenStack Platform (RHOSP) 17.1.3. This TuneD profile is the base building block to save power in RHOSP 17.1 NFV environments. For more information, see Saving power in OVS-DPDK deployments in *Configuring network functions virtualization*.

BZ#2255168

With this update, you can add load balancing capability to specific availability zones. In the **OS::Octavia::LoadBalancer** resource, use the new **availability_zone** property to specify an availability zone for the load balancer.

BZ#2255373

This enhancement updates the Block Storage (cinder) driver for Dell PowerFlex storage to support Dell PowerFlex Software version 4.5.

BZ#2261924

With this update, RHOSP 17.1 supports RHCS 7 as an external Red Hat Ceph Storage cluster.

BZ#2262266

The Shared File Systems service (manila) now includes a back-end driver to provision and manage NFS shares on Dell PowerFlex storage systems. The use of this driver is supported when the vendor publishes certification on the Ecosystem Catalog.

BZ#2262313

The Shared File Systems service (manila) now includes a back-end driver to provision and manage NFS and CIFS shares on Dell PowerStore storage systems. The use of this driver is supported when the vendor publishes certification on the Ecosystem Catalog.

BZ#2264273

This enhancement updates the Block Storage (cinder) driver for the Hewlett Packard Enterprise (HPE) 3PAR product line to support the Alletra MP Storage array.

3.1.4. Technology previews

You can test the following Technology Preview features in this release of Red Hat OpenStack Platform (RHOSP). These features provide early access to upcoming product features so that you can test

functionality and provide feedback during the development process. These features are not supported with your Red Hat subscription, and Red Hat does not recommend using them for production. For more information about the scope of support for Technology Preview features, see https://access.redhat.com/support/offerings/techpreview/.

BZ#2217663

In RHOSP 17.1, a technology preview is available for the VF-LAG transmit hash policy offload that enables load balancing at NIC hardware for offloaded traffic/flows. This hash policy is only available for layer3+4 base hashing.
To use the technology preview, verify that your templates include a bonding options parameter to enable the xmit hash policy as shown in the following example:

```
bonding_options: "mode=802.3ad miimon=100 lacp_rate=fast xmit_hash_policy=layer3+4"
```

3.1.5. Known issues

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

BZ#2163477

Currently, in RHOSP 17.1 environments that use BGP dynamic routing, the RHOSP Compute service cannot route packets sent from one of these instances to a multicast IP address destination. Therefore, instances subscribed to a multicast group fail to receive the packets sent to them. The cause is that BGP multicast routing is not properly configured on the overcloud nodes.
Workaround: Currently, there is no workaround.

BZ#2187985

Adding a load balancer member whose subnet is not in the Load-balancing service (octavia) availability zone puts the load balancer in **ERROR**. The member cannot be removed because of the **ERROR** status, making the load balancer unusable.
Workaround: Delete the load balancer.

BZ#2192913

In RHOSP environments with ML2/OVN or ML2/OVS that have DVR enabled and use VLAN tenant networks, east/west traffic between instances connected to different tenant networks is flooded to the fabric.
As a result, packets between those instances reach not only the Compute nodes where those instances run, but also any other overcloud node.

This might impact the network and it might be a security risk because the fabric sends traffic everywhere.

This bug will be fixed in a later FDP release. You do not need to perform a RHOSP update to obtain the FDP fix.

BZ#2210319

Currently, the Retbleed vulnerability mitigation in RHEL 9.2 can cause a performance drop for Open vSwitch with Data Plane Development Kit (OVS-DPDK) on Intel Skylake CPUs.
This performance regression happens only if C-states are disabled in the BIOS, Hyper-Threading Technology is enabled, and OVS-DPDK is using only one logical core of a given core.

Workaround: Assign both logical cores to OVS-DPDK or to SR-IOV guests that have DPDK running as recommended in *Configuring network functions virtualization*.

BZ#2216021

RHOSP 17.1 with the OVN mechanism driver does not support logging of flow events per port or the use of the **--target** option of the **network log create** command.
RHOSP 17.1 supports logging of flow events per security groups, using the **--resource** option of the **network log create** command. For more information, see Logging security group actions in *Configuring Red Hat OpenStack Platform networking*.

BZ#2217867

On Nvidia ConnectX-5 and ConnectX-6 NICs, when using hardware offload, some offloaded flows on a PF can cause transient performance issues on the associated VFs. This issue is specifically observed with LLDP and VRRP traffic.

BZ#2220887

The data collection service (ceilometer) reports the wrong unit for power current. Current is measured in Amperes not in Watts.

BZ#2234902

The validation **check-kernel-version** does not function correctly and reports a failure. You can ignore the failure.

BZ#2237290

The Networking service (neutron) does not prevent you from disabling or removing a networking profile, even if that profile is part of a flavor that is in use by a router. The disablement or removal of the profile can disrupt proper operation of the router.
Workaround: Before you disable or remove a networking profile, ensure that it is not part of a flavor that is currently used by a router.

BZ#2241270

The **frr-status** and **oslo-config-validator** validations report FAILED during an update. You can ignore these error messages. They are specific to the validation code and do not indicate any conditions that affect 17.1 operations. They will be fixed in a future release.

BZ#2241326

LDAP server connections are removed as expected from the Keystone LDAP pool on either **TIMEOUT** or **SERVER_DOWN** errors. The LDAP pool exhausts its connections and is unable to re-establish new ones. The **MaxConnectionReachedError** is issued. **Workaround:** Disable **LDAP pool**.

BZ#2243267

The presence of the Virtual Data Optimizer (VDO) package causes the **checkvdo** Leapp actor to fail. As a result, the Leapp upgrade fails. To complete the Leapp upgrade successfully, remove the VDO package.

BZ#2251176

The Ceph Dashboard cannot reach the Prometheus service endpoint and displays the following error message: 404 not found. This error occurs because the configuration of the VIP for the Prometheus service is not correct.
Workaround:

1. Verify haproxy is properly configured: ssh into a controller node (such as controller-0) and run **curl http://10.143.0.25:9092**. If the curl is successful, the configuration is correct.

2. If the **curl** succeeded, ssh into the controller node and update the prometheus API config in the ceph cluster:

```
$ sudo cephadm shell -- ceph dashboard set-prometheus-api-host
http://10.143.0.25:9092
```

To verify that the Ceph Dashboard can reach the Prometheus service endpoint and no longer displays the 404 not found error message, check the Ceph Dashboard UI.

BZ#2254553

Currently, in Red Hat Ceph Storage 6, **cephadm** attempts to bind the Grafana daemon to all interfaces when a valid network list is provided. This prevents the Grafana daemon from starting.

BZ#2255302

If your deployment has an external Ceph cluster with multiple file systems, you can not create a Shared File System service (Manila) share as expected.
The **cephfs_filesystem_name** driver configuration parameter that is needed to avoid this situation cannot be set using director's heat template parameters.

Workaround: Set the "cephfs_filesystem_name" parameter to specify the filesystem that the Shared File System service (Manila) must use via "ExtraConfig".

Add the parameter to an environment file as shown in the following example:

```
$ cat /home/stack/manila_cephfs_customization.yaml
parameter_defaults:
  ExtraConfig:
    manila::config::manila_config:
      cephfs/cephfs_filesystem_name:
        value: <filesystem>
```

Replace the value of <filesystem> with the appropriate name and include this environment file with the **openstack overcloud deploy** command.

BZ#2257419

The way cgroups are managed for **libvirt** changed between RHEL releases. As a result, **virsh cpu-stats** is not an officially supported part of the Red Hat OpenStack Platform (RHOSP) product. The functionality is provided by the underlying version of **libvirt** received from RHEL. Greenfield RHOSP 17.x is only supported on RHEL 9, which uses cgroups v2. The cgroups v2 API does not provide the API required to support the **virsh cpu-stats** API and this functionality is unavailable when using 17.1 on RHEL 9.
RHOSP 17.1 on RHEL 8, supported during mixed RHEL upgrades from 16.2, has been updated to make virsh cpu-stats functional while running on rhel 8.4. As a result, virsh cpu-stats functionality has been restored on RHEL 8 hosts but is unavailable on fully upgraded RHEL 9 hosts.

BZ#2259873

When upgrading RHOSP 16.2 to 17.1 on environments that use Lenovo SR650 servers, the servers fail on their first boot, displaying a blue screen stating that there is a lack of valid boot devices.
This issue is caused by Lenovo UEFI firmware resetting boot records after deployment. RHOSP director requests two changes to UEFI firmware settings. However, Lenovo hardware can only handle one request before rebooting.

Workaround: you must manually reboot the Lenovo servers to the desired operating system.

BZ#2266778

In RHOSP 17.1 environments that use the RHOSP DNS service (designate), zone transfers that involve TSIG keys can fail. The log message is: **AttributeError: 'TsigKeyring' object has no attribute 'name'**. This issue is caused by the **python3-dns** package version 2.x introducing an

incompatibility with the RHOSP DNS service. This has been fixed and will be available in a later maintenance release. **Workaround:** Currently, there is no workaround.

BZ#2267882

There is a known issue where listing the records in a zone using the RHOSP Dashboard (horizon) returns only 20 results, even if the zone contains more than 20 records. The RHOSP DNS service (designate) dashboard does not properly support pagination in the Dashboard. This has been fixed and will be available in a future RHOSP maintenance release. **Workaround:** Currently, the workaround is to use the RHOSP command line interface instead of the Dashboard.

BZ#2274468

In RHOSP 17.1 environments that use dynamic routing and the OpenStack Load-balancing service (octavia) with the OVN provider driver, there is a known issue where the Load-balancing VIP is deleted. The deletion is caused by the process that synchronizes the OVN BGP agent and the Load-balancing service.
Workaround: The workaround is to increase the reconcile interval to a very high value. Create a custom environment YAML file and add the following values:

```
parameter_defaults:
  FrrOvnBgpAgentReconcileInterval: 999999
```

For more information, see 4.11. Deploying a spine-leaf enabled overcloud .

IMPORTANT

Using this workaround means that the OVN load-balancing VIPs are operational, but the sync between the OVN BGP agent and Free Range Routing (FRR) is effectively not working. With sync non-operational, if a problem occurs during FRR configuration, FRR will not recover until the configured interval has passed.

BZ#2274663

In RHOSP environments that use dynamic routing, during a minor update, Free Range Routing (FRR) is restarted twice in a row. This occurs during updates in the following scenarios:

- from RHOSP 17.1.0 to 17.1.2 or 17.1.3.
- from RHOSP 17.1.1 to 17.1.2 or 17.1.3.
  The first restart occurs because there is a new container image. The second restart is triggered by a change to the **tripleo_frr.service** systemd file.

  These unwanted restarts were introduced in a bug fix to address BZ 2237245.

  Workaround: Perform the following steps:

IMPORTANT

This workaround requires restarting the **tripleo_frr** service and might cause network down time. Therefore, perform these steps during a maintenance window.

1. Open the config file, **/etc/systemd/system/tripleo_frr.service**.
2. After the first instance of **ExecStopPost**, add another instance of **ExecStopPost** that contains the following values:

```
ExecStopPost=/usr/bin/sleep 10
```

Example

```
[Unit]
Description=frr container
After=tripleo-container-shutdown.service
[Service]
Restart=always
ExecStart=/usr/bin/podman start frr
ExecReload=/usr/bin/podman kill --signal HUP frr
ExecStop=/usr/bin/podman stop -t 42 frr
ExecStopPost=/usr/bin/podman stop -t 42 frr
ExecStopPost=/usr/bin/sleep 10
SuccessExitStatus=137 142 143
TimeoutStopSec=84
KillMode=control-group
Type=forking
PIDFile=/run/frr.pid
[Install]
WantedBy=multi-user.target
…
```

3. Restart the **tripleo_frr** service:

```
# systemctl daemon-reload
# systemctl restart tripleo_frr
```

3.1.6. Deprecated functionality

The items in this section are either no longer supported, or will no longer be supported in a future release of Red Hat OpenStack Platform (RHOSP):

BZ#1946898

The i440FX PC machine type **pc-i440fx** was deprecated in RHEL 8. While the **pc-i440fx-*** machine types are still available, Red Hat recommends that you use the default Q35 machine type in RHOSP 17.1. Some RHOSP 17.1 features do not work with the i440FX PC machine type. For example, the VirtIO Block (**virtio-blk**) device does not work with the i440FX PC machine type in RHOSP 17.1. To use VirtIO Block as the block device for instances in RHOSP 17.1, your instances must use the Q35 machine type.

3.2. RED HAT OPENSTACK PLATFORM 17.1.2 MAINTENANCE RELEASE - JANUARY 16, 2024

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

3.2.1. Advisory list

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

RHBA-2024:0185

Red Hat OpenStack Platform 17.1.2 bug fix and enhancement advisory

RHBA-2024:0186

Updated Red Hat OpenStack Platform 17.1.2 container images

RHSA-2024:0187

Moderate: Red Hat OpenStack Platform 17.1 (python-urllib3) security update

RHSA-2024:0188

Moderate: Red Hat OpenStack Platform 17.1 (python-eventlet) security update

RHSA-2024:0189

Moderate: Red Hat OpenStack Platform 17.1 (python-werkzeug) security update

RHSA-2024:0190

Moderate: Red Hat OpenStack Platform 17.1 (GitPython) security update

RHSA-2024:0191

Moderate: Red Hat OpenStack Platform 17.1 (openstack-tripleo-common) security update

RHBA-2024:0209

Red Hat OpenStack Platform 17.1.2 bug fix and enhancement advisory

RHBA-2024:0210

Updated Red Hat OpenStack Platform 17.1.2 container images

RHBA-2024:0211

Red Hat OpenStack Platform 17.1.2 RHEL 9 director images

RHSA-2024:0212

Moderate: Red Hat OpenStack Platform 17.1 (python-django) security update

RHSA-2024:0213

Moderate: Red Hat OpenStack Platform 17.1 (python-eventlet) security update

RHSA-2024:0214

Moderate: Red Hat OpenStack Platform 17.1 (python-werkzeug) security update

RHSA-2024:0215

Moderate: Red Hat OpenStack Platform 17.1 (GitPython) security update

RHSA-2024:0216

Moderate: Red Hat OpenStack Platform 17.1 (openstack-tripleo-common) security update

RHSA-2024:0217

Moderate: Red Hat OpenStack Platform 17.1 (rabbitmq-server) security update

RHSA-2024:0263

Updated Red Hat OpenStack Platform 17.1.2 director Operator container images

3.2.2. Bug fixes

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

BZ#2108212

This update fixes an issue that disrupted connection to instances over IPv6 during migration from the OVS mechanism driver to the OVN mechanism driver.
Now you can migrate from OVS to OVN with IPv6 without experiencing the instance connection disruption.

BZ#2126725

Before this update, hard-coded certificate location operated independently of user-provided values. During deployment with custom certificate locations, services did not retrieve information from API endpoints because Transport Layer Security (TLS) verification failed. With this update, user-provided certificate locations are used during deployment.

BZ#2151219

Before this update, RHOSP director did not allow for automatically configuring nameserver (NS) records to match a parent's NS records. In RHOSP 17.1.2, this issue has been resolved by the addition of a new Orchestration service (heat) parameter, **DesignateBindNSRecords**. Administrators can use this new parameter to define the list of root NS for the domains that the DNS service (designate) populates. For more information, see *Configuring DNS as a service* .

BZ#2167428

Before this update, during a new deployment, the Identity service (keystone) was often not available during initialization of the agent-notification service. This prevented the data collection service (ceilometer) from discovering the gnocchi endpoint. As a result, metrics were not sent to gnocchi. With this update, gnocchi tries to connect to the data collection service multiple times before declaring that it cannot be reached.

BZ#2180542

This update fixes a bug that caused the **ceph-nfs** service to fail after a reboot of all Controller nodes. The Pacemaker-controlled **ceph-nfs** resource requires a runtime directory to store some process data.

Before this update, the directory was created when you installed or upgraded RHOSP. However, a reboot of the Controller nodes removed the directory, and the **ceph-nfs** service did not recover when the Controller nodes were rebooted. If all Controller nodes were rebooted, the **ceph-nfs** service failed permanently.

With this update, the directory is created before spawning the **ceph-nfs** service, and the **cephfs-nfs** service continues through reboots.

BZ#2180883

This update fixes a bug that caused **rsyslog** to stop sending logs to Elasticsearch.

BZ#2193388

Before this update, the Dashboard service (horizon) was configured to validate client TLS certificates by default, which broke the Dashboard service on all TLS everywhere (TLS-e) deployments. With this update, the Dashboard service no longer validates client TLS certificates by default, and the Dashboard service functions as expected.

BZ#2196291

This update fixes a bug that prevented non-admin users from listing or managing policy rules. Now you can allow non-admin users to list or manage policy rules.

BZ#2203785

This update fixes a permission issue that caused collectd sensubility to stop working after you rebooted a baremetal node.
Now collectd sensubility continues working after you reboot a baremetal node.

BZ#2213126

This update fixes an issue that sometimes caused the security group logging queue to stop accepting entries before reaching the limit set in **NeutronOVNLoggingRateLimit**.

You can set the maximum number of log entries per second with the parameter **NeutronOVNLoggingRateLimit**. If the log entry creation exceeds that rate, the excess is buffered in a queue up to the number of log entries that you specify in **NeutronOVNLoggingBurstLimit**.

Before this update, during short bursts, the queue sometimes stopped accepting entries before reaching the limit specified in **NeutronOVNLoggingBurstLimit**.

With this update, the **NeutronOVNLoggingBurstLimit** value affects the queue limit as expected.

BZ#2213742

This update fixes a bug that prevented TCP health monitors in UDP pools from performing as expected. Previously, the states of the pool members and the health monitors were not correctly reported. This was caused by SELinux rules that broke the use of TCP health monitors on specific port numbers in UDP pools. Now the health monitors perform correctly.

BZ#2215969

Before this update, Google Chrome did not display the list of load members correctly, which prevented members from using the dashboard to add members to a load balancer. With this update, Google Chrome displays the list of load balancer members.

BZ#2216130

Before this update, **puppet-ceilometer** did not populate the **tenant_name_discovery** parameter in the ceilometer configuration on Compute nodes. This prevented identification of the **Project name** and **User name** fields.
With this update, the addition of the **tenant_name_discovery** parameter to the Compute namespace in **puppet-ceilometer** resolves the issue. When the **tenant_name_discovery** parameter is set to **true**, the **Project name** and **User name** fields are populated.

BZ#2218596

This update fixes a bug that caused problems after a migration to the OVN mechanism driver if the original ML2/OVS environment used iptables_hybrid firewall and trunk ports.
Previously, if the original ML2/OVS environment used iptables_hybrid firewall and trunk ports, instance networking problems occurred if you recreated an instance with trunks after an event such as a hard reboot, start and stop, or node reboot.

Now you can migrate to the OVN mechanism driver if the original ML2/OVS environment uses iptables_hybrid firewall and trunk ports.

BZ#2219574

Before this update, puppet-ceilometer did not support configuring caching options for the data collection service (ceilometer). With this update, puppet-ceilometer provides configuring caching options for the data collection service (ceilometer). This support uses tripleo heat templates to provide better flexibility for configuring the caching back end.

BZ#2219613

Before this update, in RHOSP 17.1 distributed virtual router (DVR) environments, traffic was being incorrectly centralized when sent to floating IP addresses (FIPs) whose attached ports were **DOWN**. With this update, network traffic is no longer centralized if the FIP ports are in a **DOWN** state.

BZ#2220808

Before this update, the data collection service (ceilometer) did not create a resource in gnocchi because of the missing **hardware.ipmi.fan** metric in gnocchi's resource types. With this update, gnocchi reports fan metrics, which resolves the issue.

BZ#2220930

Before this update, in environments that ran the DNS service (designate), there was a known issue where the **bind9** and **unbound** services did not automatically restart if the configuration changed. With this update, the **bind9** and **unbound** services automatically restart if the configuration changes.

BZ#2222420

Before this update, in environments that used IPv6 networks that ran the RHOSP DNS service (designate), the BIND 9 backend server rejected DNS notify messages. With this update, the BIND 9 back-end server does not reject DNS notify messages.

BZ#2222825

Before this update, when you configured Nova with **[quota]count_usage_from_placement = True**, and you unshelved a shelved offloaded server, you could exceed your quota limit because a quota was not enforced. With this update, when you configure Nova with **[quota]count_usage_from_placement = True**, and you unshelve a shelved offloaded server, a quota limit is enforced.

BZ#2223294

This update fixes a bug that caused failure of the collection agent **collectd-sensubility** on RHEL 8 Compute nodes during an in-place upgrade from RHOSP 16.2 to 17.1.

BZ#2226963

Before this update, if a DCN site had 3 **DistributedComputeHCI** nodes and at least 1 **DistributedComputeHCIScaleOut** node, **cephadm** generated the incorrect spec. With this update, if a DCN site has a mix of **DistributedComputeHCI** and **DistributedComputeHCIScaleOut** nodes, **cephadm** generates the spec correctly.

BZ#2227360

Before this update, the image cache cleanup task of the NetApp NFS driver caused unpredictable slowdowns in other Block Storage services. With this update, the image cache cleanup task of the NetApp NFS driver no longer causes unpredictable slowdowns in other Block Storage services. The NetApp NFS driver also provides the **netapp_nfs_image_cache_cleanup_interval** configuration option, with a default value of 600 seconds that should be adequate for most situations.

BZ#2228818

Previously, the nova_virtlogd container did not get updated to from ubi 8 to ubi 9 as expected after a RHOSP upgrade of the Compute node to RHOSP 17.1 with RHEL 9.2. The container was updated only after rebooting the Compute node.
Now, the nova_virtlogd container gets updated to ubi 9 before the RHOSP upgrade. Note that in subsequent RHOSP updates, you must reboot the Compute node after any change to the virtlogd container, because a restart would cause workload logs to become unreachable.

BZ#2231378

Before this update, the Red Hat Ceph Storage back end of the Block Storage (cinder) backup service 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 back end. With this update, the Red Hat Ceph Storage back end 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 back end.

BZ#2232562

Before this update, **openstack overcloud deploy** did not pass the value of the **OVNAvailabilityZone** role parameter to OVS.
With this update, the **OVNAvailabilityZone** role parameter correctly passes the value as an **availability-zones** value in **external-ids:ovn-cms-options**.

The following example shows how to use the parameter in an environment file to set `OVNAvailabilityZone. Include the environment file in the deployment command.

```
ControllerParameters:
  OVNAvailabilityZone: 'az1'
```

The deployment adds **availability-zones=az1** to OVS **external-ids:ovn-cms-options**.

BZ#2233136

Before this update, when multiple values were provided in a comma-delimited list, the **CinderNetappNfsShares** parameter was incorrectly parsed. As a result, a NetApp back end with multiple NFS shares could not be defined. With this update, the **CinderNetappNfsShares** parameter is correctly parsed when provided with multiple values in a comma-delimited list. As a result, a NetApp with multiple NFS shares is correctly defined.

BZ#2233457

Before this update, the WSGI logs for the **cinder-api** service were not stored in a persistent location, which caused you to not be able to view the logs to troubleshoot issues. With this update, the WSGI logs are stored on the controller nodes where the **cinder-api** service runs in the **/var/log/containers/httpd/cinder-api** directory, which resolves the issue.

BZ#2233487

Before this update, if you used RHOSP dynamic routing in your RHOSP environment and you created a load balancer by using the RHOSP Load-balancing service (octavia), the latency between the Controller nodes might have caused the OVN provider driver to fail. With this update, load balancers are successfully created when using the OVN provider driver on Controller nodes that are experiencing latency.

BZ#2235621

Before this update, the RHOSP upgrade from 16.2 to 17.1 failed when pulling images from **registry.redhat.io** because the upgrade playbook did not include the Podman registry login task. This issue is resolved in RHOSP 17.1.2.

BZ#2237245

With this update, RHOSP 17.1 environments that use dynamic routing that are updating to RHOSP 17.1.2 now work properly. RHOSP director now successfully updates the Free Range Routing (FRR) component without requiring any workaround.

BZ#2237251

Before this update, RHOSP environments that used the Load-balancing service (octavia) with the OVN provider and a health monitor caused the load-balancing pool to display a fake member's status as **ONLINE**. With this update, if you use a health monitor for a pool, the fake load-balancing pool member now has the **ERROR** operating status and the Load Balancer/Listener/Pool operating statuses are updated accordingly.

BZ#2237866

Before this update, configuring caching parameters for ceilometer were not supported. With this update, for caching, ceilometer uses the **dogpile.cache.memcached** back end. If you manually disable caching, celiometer uses the **oslo_cache.dict** back end.

BZ#2240591

Before this update, calling the member batch update API triggered race conditions in the Octavia API service, which caused the load balancer to be stuck in the "PENDING_UPDATE" provisioning_status. With this update, calling the member batch update API does not trigger race conditions, which resolves the issue.

BZ#2242605

Before this update, an upgrade from RHOSP 16.2 to 17.1 failed on environments that were not connected to the internet because the **infra_image** value was not defined. The **overcloud_upgrade_prepare.sh** script tried to pull **registry.access.redhat.com/ubi8/pause** instead, which caused an error. The issue is resolved in RHOSP 17.1.2.

BZ#2244631

Before this update, performing a manual OVN DB sync while the OVN metadata and the OVN LB health monitor ports were present in the same environment caused the OVN DB sync to delete one of the ports. If the OVN metadata port was deleted, you lost communication with the VMs. With this update, a manual OVN DB sync does not delete one of the ports because the OVN-provider uses the **ovn-lb-hm:distributed** value for the **device_owner** parameter. OVN provider updates existing OVN LB Health Monitor ports to the **ovn-lb-hm:distributed** value.

BZ#2246563

Before this update, Director did not include the puppet modules and the heat templates that you needed to configure the Pure Flashblade driver with your Red Hat Openstack Shared File System Service (manila). With this update, Director now includes the necessary puppet modules and heat templates for your configuration.

3.2.3. Enhancements

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

BZ#1759007

The upgrade of multi-cell environments is now supported.

BZ#1813561

With this update, the Load-balancing service (octavia) supports HTTP/2 load balancing by using the Application Layer Protocol Negotiation (ALPN) for listeners and pools that are enabled with Transport Layer Security (TLS). The HTTP/2 protocol improves performance by loading pages faster.

BZ#1816766

This enhancement adds support for uploading compressed images to the Image service (glance). You can use the image decompression plugin to optimize network bandwidth by reducing image upload times and storage consumption on hosts.

BZ#2222699

This update fixes a bug that set the wrong MTU value on tenant networks that were changed from VXLAN to Geneve after a migration from the OVS mechanism driver to the OVN mechanism driver. Before this update, the **cloud-init** package overrode the value that was correctly set by the DHCP server.
For example, after a migration from the OVS mechanism driver with VXLAN to the OVN mechanism driver to Geneve with a 1442 MTU, cloud-init reset the MTU to 1500.

With this update, the value set by the DHCP server persists.

BZ#2233695

This enhancement adds support for the Revert to Snapshot feature for iSCSI, FC, and NFS drivers with FlexVol pool. **Limitations:** This feature does not support FlexGroups. Also, you can revert to only the most recent snapshot of a Block Storage volume.

BZ#2237500

This update clarifies an error message produced by **openstack-tripleo-validations**. Previously, if a host was not found when you ran a validation, the command reported the status as **FAILED**. Now the Status is reported as **SKIPPED**.

3.2.4. Technology previews

You can test the following Technology Preview features in this release of Red Hat OpenStack Platform (RHOSP). These features provide early access to upcoming product features so that you can test

functionality and provide feedback during the development process. These features are not supported with your Red Hat subscription, and Red Hat does not recommend using them for production. For more information about the scope of support for Technology Preview features, see https://access.redhat.com/support/offerings/techpreview/.

BZ#1848407

In RHOSP 17.1, a technology preview is available for the Stream Control Transmission Protocol (SCTP) in the Load-balancing service (octavia). Users can create SCTP listeners and attach SCTP pools in a load balancer.

BZ#2217663

In RHOSP 17.1, a technology preview is available for the VF-LAG transmit hash policy offload that enables load balancing at NIC hardware for offloaded traffic/flows. This hash policy is only available for layer3+4 base hashing.
To use the technology preview, verify that your templates include a bonding options parameter to enable the xmit hash policy as shown in the following example:

```
bonding_options: "mode=802.3ad miimon=100 lacp_rate=fast xmit_hash_policy=layer3+4"
```

3.2.5. Known issues

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

BZ#2034801

A RHOSP deployment can fail when a very large number of virtual functions (VFs) are created per physical function (PF). The NetworkManager issues a DHCP request on all of them, leading to failures in the NetworkManager service.
For example, this issue occurred during a deployment that included 256 VFs across 4 PFs.

Workaround: Avoid creating a very large number of VFs per PF.

BZ#2107599

Do not change **binding:vnic_type** on a port that is attached to an instance. Doing so causes **nova_compute** to go into a restart loop if it is restarted.

BZ#2160481

In RHOSP 17.1 environments that use BGP dynamic routing, there is currently a known issue where floating IP (FIP) port forwarding fails.
When FIP port forwarding is configured, packets sent to a specific destination port with a destination IP that equals the FIP are redirected to an internal IP from a RHOSP Networking service (neutron) port. This occurs regardless of the protocol that is used: TCP, UDP, and so on.

When BGP dynamic routing is configured, the routes to the FIPs used to perform FIP port forwarding are not exposed, and these packets cannot reach their final destinations.

Workaround: Currently, there is no workaround.

BZ#2163477

In RHOSP 17.1 environments that use BGP dynamic routing, there is currently a known issue affecting instances connected to provider networks. The RHOSP Compute service cannot route packets sent from one of these instances to a multicast IP address destination. Therefore, instances subscribed to a multicast group fail to receive the packets sent to them. The cause is that BGP multicast routing is not properly configured on the overcloud nodes. **Workaround:** Currently, there is no workaround.

BZ#2178500

If a volume refresh fails when using the **nova-manage** CLI, this causes the instance to stay in a locked state.

BZ#2187985

Adding a load balancer member whose subnet is not in the Load-balancing service (octavia) availability zone puts the load balancer in **ERROR**. The member cannot be removed because of the **ERROR** status, making the load balancer unusable.
Workaround: Delete the load balancer.

BZ#2192913

In RHOSP environments with ML2/OVN or ML2/OVS that have DVR enabled and use VLAN tenant networks, east/west traffic between instances connected to different tenant networks is flooded to the fabric.
As a result, packets between those instances reach not only the Compute nodes where those instances run, but also any other overcloud node.

This might impact the network and it might be a security risk because the fabric sends traffic everywhere.

This bug will be fixed in a later FDP release. You do not need to perform a RHOSP update to obtain the FDP fix.

BZ#2210319

Currently, the Retbleed vulnerability mitigation in RHEL 9.2 can cause a performance drop for Open vSwitch with Data Plane Development Kit (OVS-DPDK) on Intel Skylake CPUs.
This performance regression happens only if C-states are disabled in the BIOS, Hyper-Threading Technology is enabled, and OVS-DPDK is using only one logical core of a given core.

Workaround: Assign both logical cores to OVS-DPDK or to SRIOV guests that have DPDK running as recommended in the NFV configuration guide.

BZ#2216021

RHOSP 17.1 with the OVN mechanism driver does not support logging of flow events per port or the use of the **--target** option of the **network log create** command.
RHOSP 17.1 supports logging of flow events per security groups, using the **--resource** option of the **network log create** command. For more information, see Logging security group actions in *Configuring Red Hat OpenStack Platform networking*.

BZ#2217867

On Nvidia ConnectX-5 and ConnectX-6 NICs, when using hardware offload, where some offloaded flows on a PF can cause transient performance issues on the associated VFs. This issue is specifically observed with LLDP and VRRP traffic.

BZ#2220887

The data collection service (ceilometer) does not filter separate power and current metrics.

BZ#2222683

Currently, there is no support for Multi-RHEL for the following deployment architectures:

- Edge (DCN)
- ShiftOnStack

- Director operator-based deployments
  Workaround: Use only a single version of RHEL across your RHOSP deployment when operating one of the listed architectures.

BZ#2223916

In RHOSP 17.1 GA environments that use the ML2/OVN mechanism driver, floating IP port forwarding does not function correctly.
FIP port forwarding should be centralized on the Controller or the Networker nodes. Instead, VLAN and flat networks distribute north-south network traffic when FIPs are used.

Workaround: To resolve this problem and force FIP port forwarding through the centralized gateway node, either set the RHOSP Orchestration service (heat) parameter **NeutronEnableDVR** to **false**, or use Geneve instead of VLAN or flat project networks.

BZ#2224236

In this release of RHOSP, 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#2231893

The metadata service can become unavailable after the metadata agent fails in multiple attempts to start a malfunctioning HAProxy child container. The metadata agent logs an error message similar to: `ProcessExecutionError: Exit code: 125; Stdin: ; Stdout: Starting a new child container neutron-haproxy-ovnmeta-<uuid>".
Workaround: Run **podman kill <_container name_>** to stop the problematic haproxy child container.

BZ#2231960

When a Block Storage volume uses the Red Hat Ceph Storage back end, a volume cannot be removed when a snapshot is created from this volume and then a volume clone is created from this snapshot. In this case, you cannot remove the original volume while the volume clone exists.

BZ#2237290

The Networking service (neutron) does not prevent you from disabling or removing a networking profile, even if that profile is part of a flavor that is in use by a router. The disablement or removal of the profile can disrupt proper operation of the router.
Workaround: Before you disable or remove a networking profile, ensure that it is not part of a flavor that is currently used by a router.

BZ#2241270

The **frr-status** and **oslo-config-validator** validations report FAILED during an update. You can ignore these error messages. They are specific to the validation code and do not indicate any conditions that affect 17.1 operations. They will be fixed in a future release.

BZ#2241326

LDAP server connections are removed as expected from the Keystone LDAP pool on either **TIMEOUT** or **SERVER_DOWN** errors. The LDAP pool exhausts its connections and is unable to re-establish new ones. The **MaxConnectionReachedError** is issued. **Workaround:** Disable **LDAP pool**.

BZ#2242439

With **localnet_learn_fdb** enabled, packet loss can occur in traffic between two instances hosted by different Compute nodes. This is a core OVN issue. To avoid the issue, do not enable **localnet_learn_fdb**.

BZ#2249690

If there are multiple clusters in DCN FFU, Ceph cluster upgrades fail because they cannot find the **ceph-ansible** package as it is removed during the first Ceph cluster upgrade.

BZ#2251176

The Ceph Dashboard cannot reach the Prometheus service endpoint and displays the following error message: 404 not found. This error occurs because the configuration of the VIP for the Prometheus service is not correct.
Workaround:

1. Verify haproxy is properly configured: ssh into a controller node (such as controller-0) and run **curl http://10.143.0.25:9092**. If the curl is successful, the configuration is correct.

2. If the **curl** succeeded, ssh into the controller node and update the prometheus API config in the ceph cluster:

```
$ sudo cephadm shell -- ceph dashboard set-prometheus-api-host
http://10.143.0.25:9092
```

   To verify that the Ceph Dashboard can reach the Prometheus service endpoint and no longer displays the 404 not found error message, check the Ceph Dashboard UI.

BZ#2252723

Some AMD environments fail to boot when provisioned with the overcloud-hardened-uefi-full.raw image, due to the included kernel argument **console=ttyS0**. As a result, the boot sequence halts with no diagnostic nor error message.
Workaround: Run the following commands to edit the overcloud image:

```
sudo yum install guestfs-tools -y

sudo systemctl start libvirtd

sudo virt-customize -a /var/lib/ironic/images/overcloud-hardened-uefi-full.raw \
  --run-command "sed -i 's/console=ttyS0 //g' /etc/default/grub" \
  --run-command "grub2-mkconfig -o /boot/grub2/grub.cfg" \
  --run-command "grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg"
```

After running those commands, you can provision the AMD nodes using the provision command.

BZ#2254036

During director-deployed Ceph upgrade, if the CephClusterName variable was overridden to a value other than "ceph", then the upgrade process fails. All distributed compute nodes (DCN) deployments override this variable.

BZ#2254553

In Red Hat Ceph Storage 6, there is currently a known issue where **cephadm** attempts to bind the Grafana daemon to all interfaces when a valid network list is provided. This prevents the Grafana daemon from starting.

BZ#2254994

In RHOSP 17.1.2 environments that contain Load-balancing service (octavia) health monitor ports from a previous version, running **neutron-db-sync-tool** might randomly delete any of those pre-existing ports or OVN metadata ports. This unintended port deletion, results in a loss of health monitor capacity, or communication loss with the affected instances.
Workaround: Manually update the 'device_owner' field on existing Load-balancing service health monitor ports to the value of **ovn-lb-hm:distributed**. Doing so ensures that if the **neutron-db-sync-tool** is launched, the health monitor or OVN metadata ports are not adversely impacted.

BZ#2255302

If your deployment has an external Ceph cluster with multiple file systems, you can not create a Shared File System service (Manila) share as expected.
The **cephfs_filesystem_name** driver configuration parameter that is needed to avoid this situation cannot be set using director's heat template parameters.

Workaround: Set the "cephfs_filesystem_name" parameter to specify the filesystem that the Shared File System service (Manila) must use via "ExtraConfig".

Add the parameter to an environment file as shown in the following example:

```
$ cat /home/stack/manila_cephfs_customization.yaml
parameter_defaults:
  ExtraConfig:
    manila::config::manila_config:
      cephfs/cephfs_filesystem_name:
        value: <filesystem>
```

Replace the value of <filesystem> with the appropriate name and include this environment file with the **openstack overcloud deploy** command.

BZ#2255324

A director bug can disrupt or crash client workloads during updates or upgrades to any RHOSP 17.1 version. This bug affects deployments that enable the RHOSP Shared File Systems service (manila) with the CephFS-via-NFS backend.
The bug causes deletion of Ceph NFS export information during update or upgrade operations. This export information is created by the Shared File System Service (manila) when users set up "access rules" on their shares.

When the NFS server goes into a recovery mode, Client workloads can hang and eventually crash if they were actively reading or writing to NFS shares.

Workaround: See Manila shares with Red Hat OpenStack 17.1 can be abruptly disconnected due to export information loss.

3.3. RED HAT OPENSTACK PLATFORM 17.1.1 MAINTENANCE RELEASE - SEPTEMBER 20, 2023

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

3.3.1. Advisory list

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

RHBA-2023:5134

Release of containers for OSP 17.1

RHBA-2023:5135

Release of components for OSP 17.1

RHBA-2023:5136

Release of containers for OSP 17.1

RHBA-2023:5137

Red Hat OpenStack Platform 17.1 RHEL 9 deployment images

RHBA-2023:5138

Release of components for OSP 17.1

3.3.2. Bug fixes

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

BZ#2184834

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 multi-attach volume had been deprecated for removal because it is unsafe and 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 only supports creating a multi-attach volume by using a multi-attach volume-type. Therefore some Block Storage API requests that used to work will be rejected with a 400 (Bad Request) response code and an informative error message.

BZ#2222543

This update fixes a bug that negatively affected OVN database operation after the replacement of a bootstrap Controller node. Before this update, you could not use the original bootstrap Controller node hostname and IP address for the replacement Controller node, because the name reuse caused issues with OVN database RAFT clusters.
Now, you can use the original hostname and IP address for the replacement Controller node.

BZ#2222589

Before this update, during the upgrade from RHOSP 16.2 to 17.1, the director upgrade script stopped executing when upgrading Red Hat Ceph Storage 4 to 5 in a director-deployed Ceph Storage environment that used IPv6. This issue is resolved in RHOSP 17.1.1.

BZ#2224527

Before this update, the upgrade procedure from RHOSP 16.2 to 17.1 failed when RADOS Gateway (RGW) was deployed as part of director-deployed Red Hat Ceph Storage because HAProxy did not restart on the next stack update. This issue was resolved in Red Hat Ceph Storage 5.3.5 and no longer impacts RHOSP upgrades.

BZ#2226366

Before this update, when retyping **in-use** Red Hat Ceph Storage (RHCS) volumes to store the volume in a different pool than its current location, data could be corrupted or lost. With this update, the Block Storage RHCS back end resolves this issue.

BZ#2227199

Before this update, in RHOSP 17.1 environments that used the Load-balancing service (octavia) with the OVN service provider driver, load balancer health checks for floating IP addresses (FIPs) were not properly populated with the protocol port. Requests to the FIPs were incorrectly distributed to load balancer members that were in the `ERROR`state.
With this update, the issue is resolved, and any new load balancer health checks for floating IP addresses (FIPs) are properly populated with the protocol port. If you created health monitors before deploying this update, you must recreate them to resolve the port issue.

BZ#2229750

Before this update, when specifying an availability zone (AZ) when creating a Block Storage volume backup, the AZ was ignored, which could cause the backup to fail. With this update, the Block Storage backup service resolves this issue.

BZ#2229761

Before this update, a race condition in the deployment steps for **ovn_controller** and **ovn_dbs** caused **ovn_dbs** to be upgraded before **ovn_controller. If `ovn_controller** is not upgraded before **ovn_dbs**, an error before the restart to the new version causes packet loss. In RHOSP 17.1.1, this issue has been resolved.

BZ#2229767

Before this update, when you upgraded Red Hat Ceph Storage 4 to 5 during the upgrade from RHOSP 16.2 to 17.1, the overcloud upgrade failed because the containers that were associated with **ceph-nfs-pacemaker** were down, impacting the Shared File Systems service (manila). This issue is resolved in RHOSP 17.1.1.

3.3.3. Enhancements

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

BZ#2210151

In RHOSP 17.1.1, the RHOSP Orchestration service (heat) parameter, **FrrBgpAsn**, can now be set on a per-role basis instead of being a global parameter for RHOSP 17.1 environments that use RHOSP dynamic routing.

BZ#2229026

In RHOSP 17.1.1, the **tripleo_frr_bgp_peers** role-specific parameter can now be used to specify a list of IP addresses or hostnames for Free Range Routing (FRR) to peer with.
Example

```
ControllerRack1ExtraGroupVars:
  tripleo_frr_bgp_peers: ["172.16.0.1", "172.16.0.2"]
```

3.3.4. Technology previews

The items listed in this section are provided as Technology Previews in this release of 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#1813561

With this update, the Load-balancing service (octavia) supports HTTP/2 load balancing by using the Application Layer Protocol Negotiation (ALPN) for listeners and pools that are enabled with Transport Layer Security (TLS). The HTTP/2 protocol improves performance by loading pages

faster.

BZ#1848407

In RHOSP 17.1, a technology preview is available for the Stream Control Transmission Protocol (SCTP) in the Load-balancing service (octavia). Users can create SCTP listeners and attach SCTP pools in a load balancer.

BZ#2211796

This release includes a Technology Preview of the optional feature that you can use to define custom router flavors and create routers with the custom router flavors.
For more information, see Creating custom virtual routers with router flavors.

BZ#2217663

In RHOSP 17.1, a technology preview is available for the VF-LAG transmit hash policy offload that enables load balancing at NIC hardware for offloaded traffic/flows. This hash policy is only available for layer3+4 base hashing.

3.3.5. Known issues

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

BZ#2108212

If you use IPv6 to connect to instances during migration to the OVN mechanism driver, connection to the instances might be disrupted for up to several minutes when the ML2/OVS services are stopped. The router advertisement daemon **radvd** for IPv6 is stopped during migration to the OVN mechanism driver. While **radvd** is stopped, router advertisements are no longer broadcast. This broadcast interruption results in instance connection loss over IPv6. IPv6 communication is automatically restored once the new ML2/OVN services start.

Workaround: To avoid the potential disruption, use IPv4 instead.

BZ#2126725

Hard-coded certificate location operates independently of user-provided values. During deployment with custom certificate locations, services do not retrieve information from API endpoints because Transport Layer Security (TLS) verification fails.

BZ#2144492

If you migrate a RHOSP 17.1.0 ML2/OVS deployment with distributed virtual routing (DVR) to ML2/OVN, the floating IP (FIP) downtime that occurs during ML2/OVN migration can exceed 60 seconds.

BZ#2151290

In RHOSP 17.1.1, director does not allow for automatically configuring NS records to match a parent's NS records. Workaround: Until an automated workaround is provided in a future release, administrators can manually change the Orchestration service (heat) template file that resides on the undercloud in **/usr/share/ansible/roles/designate_bind_pool/templates/**. In the Jinja template, **pools.yaml.j2**, remove the code following the line containing **ns_records** until the next empty line (lines 13-16) and insert appropriate values for their infrastructure. Finally, administrators should redeploy the overcloud.
Example

```
ns_records:
  - hostname: ns1.desiexample.com
    priority: 1
```

```
- hostname: ns2.desiexample.com
  priority: 2
```

BZ#2160481

In RHOSP 17.1 environments that use BGP dynamic routing, there is currently a known issue where floating IP (FIP) port forwarding fails.
When FIP port forwarding is configured, packets sent to a specific destination port with a destination IP that equals the FIP are redirected to an internal IP from a RHOSP Networking service (neutron) port. This occurs regardless of the protocol that is used: TCP, UDP, and so on.

When BGP dynamic routing is configured, the routes to the FIPs used to perform FIP port forwarding are not exposed, and these packets cannot reach their final destinations.

Currently, there is no workaround.

BZ#2163477

In RHOSP 17.1 environments that use BGP dynamic routing, there is currently a known issue affecting instances connected to provider networks. The RHOSP Compute service cannot route packets sent from one of these instances to a multicast IP address destination. Therefore, instances subscribed to a multicast group fail to receive the packets sent to them. The cause is that BGP multicast routing is not properly configured on the overcloud nodes. Currently, there is no workaround.

BZ#2167428

In RHOSP 17.1.1, there is a known issue during a new deployment where the RHOSP Identity service (keystone) is often not available when the **agent-notification** service is initializing. This prevents ceilometer from discovering the gnocchi endpoint. As a result, metrics are not sent to gnocchi.
Workaround: Restart the agent-notification service on the Controller node:

```
$ sudo systemctl restart tripleo_ceilometer_agent_notification.service
```

BZ#2178500

If a volume refresh fails when using the nova-manage CLI, this causes the instance to stay in a locked state.

BZ#2180542

The Pacemaker-controlled **ceph-nfs** resource requires a runtime directory to store some process data. The directory is created when you install or upgrade RHOSP. Currently, a reboot of the Controller nodes removes the directory, and the **ceph-nfs** service does not recover when the Controller nodes are rebooted. If all Controller nodes are rebooted, the **ceph-nfs** service fails permanently.
Workaround: If you reboot a Controller node, log into the Controller node and create a **/var/run/ceph** directory:

$ mkdir -p /var/run/ceph

Repeat this step on all Controller nodes that have been rebooted. If the **ceph-nfs-pacemaker** service has been marked as failed, after creating the directory, execute the following command from any of the Controller nodes:

$ pcs resource cleanup

BZ#2180883

Currently, rsyslog stops sending logs to Elasticsearch when Logrotate archives all log files once a day.

Workaround: Add "RsyslogReopenOnTruncate: true" to your environment file during deployment so that Rsyslog reopens all log files on log rotation.

Currently, RHOSP 17.1 is shipped with puppet-rsyslog module, which causes Director to configure rsyslog incorrectly.

Workaround: Manually apply patch [1] in **/usr/share/openstack-tripleo-heat-templates/deployment/logging/rsyslog-container-puppet.yaml** before deployment to configure Rsyslog correctly.

[1] https://github.com/openstack/tripleo-heat-templates/commit/ce0e3a9a94a4fce84dd70b6098867db1c86477fb

BZ#2192913

In RHOSP environments with ML2/OVN or ML2/OVS that have DVR enabled and use VLAN tenant networks, east/west traffic between instances connected to different tenant networks is flooded to the fabric.
As a result, packets between those instances reach not only the Compute nodes where those instances run, but also any other overcloud node.

This could cause an impact on the network and it could be a security risk because the fabric sends traffic everywhere.

This bug will be fixed in a later FDP release. You do not need to perform a RHOSP update to obtain the FDP fix.

BZ#2196291

Currently, custom SRBAC rules do not permit list policy rules to non-admin users. As a consequence, non-admin users can not list or manage these rules. Current workarounds include either disabling SRBAC, or modifying the SRBAC custom rule to permit this action.

BZ#2203785

Currently, there is a permission issue that causes collectd sensubility to stop working after you reboot a baremetal node. As a consequence, sensubility stops reporting container health.
Workaround: After rebooting an overcloud node, manually run the following command on the node:
sudo podman exec -it collectd setfacl -R -m u:collectd:rwx /run/podman

BZ#2210319

Currently, the Retbleed vulnerability mitigation in RHEL 9.2 can cause a performance drop for Open vSwitch with Data Plane Development Kit (OVS-DPDK) on Intel Skylake CPUs.
This performance regression happens only if C-states are disabled in the BIOS, hyper-threading is enabled, and OVS-DPDK is using only one hyper-thread of a given core.

Workaround: Assign both hyper-threads of a core to OVS-DPDK or to SRIOV guests that have DPDK running as recommended in the NFV configuration guide.

BZ#2210873

In RHOSP 17.1.1 Red Hat Ceph Storage (RHCS) environments, setting crush rules fail with an **assimilate.conf not found** error. This problem will be fixed in a later RHOSP release.

BZ#2213126

The logging queue that buffers excess security group log entries sometimes stops accepting entries before the specified limit is reached. As a workaround, you can set the queue length higher than the number of entries you want it to hold.

You can set the maximum number of log entries per second with the parameter **NeutronOVNLoggingRateLimit**. If the log entry creation exceeds that rate, the excess is buffered in a queue up to the number of log entries that you specify in **NeutronOVNLoggingBurstLimit**.

The issue is especially evident in the first second of a burst. In longer bursts, such as 60 seconds, the rate limit is more influential and compensates for burst limit inaccuracy. Thus, the issue has the greatest proportional effect in short bursts.

Workaround: Set **NeutronOVNLoggingBurstLimit** at a higher value than the target value. Observe and adjust as needed.

BZ#2213742

TCP health monitors in UDP pools might not work as expected, depending on the port number that is used by the monitor. Also the status of the pool members and the health monitors are not correct. This is caused by SELinux rules that break the use of TCP health monitors on specific port numbers in UDP pools.
Workaround (if any): Currently, there is no workaround.

BZ#2216021

RHOSP 17.1 with the OVN mechanism driver does not support logging of flow events per port or the use of the **--target** option of the **network log create** command.
RHOSP 17.1 supports logging of flow events per security groups, using the **--resource** option of the **network log create** command. See "Logging security group actions" in *Networking with RHOSP*.

BZ#2216130

Currently, **puppet-ceilometer** does not populate the **tenant_name_discovery** parameter in the data collection service (ceilometer) configuration on Compute nodes. This causes the **Project name** and **User name** fields to not be identified. Currently, there is no workaround for this issue.

BZ#2217867

There is currently a known issue on Nvidia ConnectX-5 and ConnectX-6 NICs, when using hardware offload, where some offloaded flows on a PF can cause transient performance issues on the associated VFs. This issue is specifically observed with LLDP and VRRP traffic.

BZ#2218596

Do not migrate to the OVN mechanism driver if your original ML2/OVS environment uses iptables_hybrid firewall and trunk ports. In the migrated environment, instance networking problems occur if you recreate an instance with trunks after an event such as a hard reboot, start and stop, or node reboot. As a workaround, you can switch from the iptables hybrid firewall to the OVS firewall before migrating.

BZ#2219574

The data collection service (ceilometer) does not provide a default caching back end, which can cause some services to be overloaded when polling for metrics.

BZ#2219603

In RHOSP 17.1 GA, the DNS service (designate) is misconfigured when secure role-based access control (sRBAC) is enabled. The current sRBAC policies contain incorrect rules for designate and must be corrected for designate to function correctly. A possible workaround is to apply the following patch on the undercloud server and redeploy the overcloud:
https://review.opendev.org/c/openstack/tripleo-heat-templates/+/888159

BZ#2219613

In RHOSP 17.1 distributed virtual router (DVR) environments, the **external_mac** variable is improperly being removed for ports in the **DOWN** status which results in centralized traffic for short periods.

BZ#2219830

In RHOSP 17.1, there is a known issue of transient packet loss where hardware interrupt requests (IRQs) are causing non-voluntary context switches on OVS-DPDK PMD threads or in guests running DPDK applications.
This issue is the result of provisioning large numbers of VFs during deployment. VFs need IRQs, each of which must be bound to a physical CPU. When there are not enough housekeeping CPUs to handle the capacity of IRQs, **irqbalance** fails to bind all of them and the IRQs overspill on isolated CPUs.

Workaround: You can try one or more of these actions:

- Reduce the number of provisioned VFs to avoid unused VFs remaining bound to their default Linux driver.
- Increase the number of housekeeping CPUs to handle all IRQs.
- Force unused VF network interfaces down to avoid IRQs from interrupting isolated CPUs.
- Disable multicast and broadcast traffic on unused, down VF network interfaces to avoid IRQs from interrupting isolated CPUs.

BZ#2220808

In RHOSP 17.1, there is a known issue where the data collection service (ceilometer) does not report airflow metrics. This problem is caused because the data collection service is missing a gnocchi resource type, **hardware.ipmi.fan**. Currently, there is no workaround.

BZ#2220887

The data collection service (ceilometer) does not filter separate power and current metrics.

BZ#2220930

In RHOSP 17.1 that run the DNS service (designate), there is a known issue where the **bind9** and **unbound** services are not restarted if the configuration changes.
Workaround: Manually restart the containers by running the following commands on each controller:

```
$ sudo systemctl restart tripleo_designate_backend_bind9
$ sudo systemctl restart tripleo_unbound
```

BZ#2222420

In RHOSP 17.1.1 environments that use IPv6 networks that run the RHOSP DNS service (designate), the BIND 9 back end server can reject DNS notify messages. This issue is caused because there are often multiple IP addresses for the same network on the same interface, and it can appear that the messages are emanating from sources other than the designate Worker services.
Workaround: Apply the following patches:

- https://review.opendev.org/c/openstack/tripleo-ansible/+/888300
- https://review.opendev.org/c/openstack/tripleo-heat-templates/+/888786
  After you apply the patches, manually restart the configuration in the BIND 9 servers by running:

```
$ sudo systemctl restart tripleo_designate_backend_bind9
```

BZ#2222683

Currently, there is no support for Multi-RHEL for the following deployment architectures:

- Edge (DCN)
- ShiftOnStack
- Director operator-based deployments
  Workaround: Use only a single version of RHEL across your RHOSP deployment when operating one of the listed architectures.

BZ#2223294

There is a known issue when performing an in-place upgrade from RHOSP 16.2 to 17.1 GA. The collection agent, **collectd-sensubility** fails to run on RHEL 8 Compute nodes.
Workaround: On affected nodes edit the file, **/var/lib/container-config-scripts/collectd_check_health.py**, and replace **"healthy: .State.Health.Status}"** with **"healthy: .State.Healthcheck.Status}"**/ on line 26.

BZ#2223916

In RHOSP 17.1 GA environments that use the ML2/OVN mechanism driver, there is a known issue with floating IP port forwarding not working correctly. This problem is caused because VLAN and flat networks distribute north-south network traffic when FIPs are used, and, instead, FIP port forwarding should be centralized on the Controller or the Networker nodes.
Workaround: To resolve this problem and force FIP port forwarding through the centralized gateway node, either set the RHOSP Orchestration service (heat) parameter **NeutronEnableDVR** to **false**, or use Geneve instead of VLAN or flat project networks.

BZ#2224236

In this release of 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#2225205

Outdated upgrade orchestration logic overrides the existing Pacemaker authkey during the Fast Forward Upgrade (FFU) procedure, preventing Pacemaker from connecting to **pacemaker_remote** running on Compute nodes when Instance HA is enabled. As a result, the upgrade fails and **pacemaker_remote** running on Compute nodes is unreachable from the central cluster. Contact Red Hat support to receive instructions on how to perform FFU if Instance HA is configured.

BZ#2227360

The image cache cleanup task of the NetApp NFS driver can cause unpredictable slowdowns in other Block Storage services. There is currently no workaround for this issue.

BZ#2229937

When **collectd sensubility** fails to create a sender, it does not close the link to the sender. Long-

running open links that fail can cause issues in the bus, which cause **collectd sensubility** to stop working. Workaround: Restart the **collectd** container on affected overcloud nodes to recover **collectd sensubility**.

BZ#2231378

If you choose Red Hat Ceph Storage as the back end for your Block Storage (cinder) backup service repository, then you can only restore backed up volumes to a RBD-based Block Storage back end. There is currently no workaround for this.

BZ#2231893

The metadata service can become unavailable after the metadata agent fails in multiple attempts to start a malfunctioning HAProxy child container. The metadata agent logs an error message similar to: `ProcessExecutionError: Exit code: 125; Stdin: ; Stdout: Starting a new child container neutron-haproxy-ovnmeta-<uuid>".
Workaround: Run **podman kill <_container name_>** to stop the problematic haproxy child container.

BZ#2231960

When a Block Storage volume uses the Red Hat Ceph Storage back end, a volume cannot be removed when a snapshot is created from this volume and then a volume clone is created from this snapshot. In this case, you cannot remove the original volume while the volume clone exists.

BZ#2232562

The **OVNAvailabilityZone Role** parameter is not recognized as expected, which causes availability zone configuration to fail in OVN.
Workaround: Use the **OVNCMSOptions** parameter to configure OVN availability zones. For example:

```
ControllerParameters:
  OVNCMSOptions: 'enable-chassis-as-gw,availability-zones=az1'
```

BZ#2233487

In RHOSP 17.1 GA environments that use RHOSP dynamic routing, there is a known issue where creating a load balancer using the RHOSP Load-balancing service with the OVN provider driver might fail. This failure can occur when there is latency between controller nodes. There is no workaround.

BZ#2235621

The RHOSP upgrade from 16.2 to 17.1 fails when pulling images from **registry.redhat.io** because the upgrade playbook does not include the podman registry login task. Contact your Red Hat support representative for a hotfix. A fix is expected in a later RHOSP release.

BZ#2237245

In RHOSP 17.1 environments that use dynamic routing, updating to RHOSP 17.1.1 does not work properly. Specifically, Free Range Routing (FRR) components are not updated.
Workaround: Apply the following patches on the undercloud before updating RHOSP 17.1:

- https://review.opendev.org/c/openstack/tripleo-ansible/+/893486
- https://review.opendev.org/c/openstack/tripleo-ansible/+/894171
- https://review.opendev.org/c/openstack/tripleo-heat-templates/+/893616

BZ#2237251

In RHOSP 17.1.1 environments that use the Load-balancing service (octavia) with the OVN provider driver with a health monitor, the pool load-balancing status incorrectly displays fake members as **ONLINE**. If no health monitor is being used, then the status fake member displays a normal operation

of **NO_MONITOR**.
Fake load-balancing pool members can occur when a member is not valid, such as when there is a typographical error in the member's IP address. Health monitors configured for the pool perform no health checks on the fake member, and the global operating status incorrectly considers the fake member as **ONLINE** when it calculates the pool's status. Furthermore, if all other members in a pool are in **ERROR** operating status, an incorrect **DEGRADED** operating status is assigned to the pool instead of **ERROR** because a member of the pool is a fake member with an incorrect **ONLINE** status.

Workaround: Currently, there are no workarounds for this issue.

BZ#2237290

The Networking service (neutron) does not prevent you from disabling or removing a networking profile, even if that profile is part of a flavor that is in use by a router. The disablement or removal of the profile can disrupt proper operation of the router.
Workaround: Before you disable or remove a networking profile, ensure that it is not part of a flavor that is currently used by a router.

3.4. RED HAT OPENSTACK PLATFORM 17.1 GA - AUGUST 17, 2023

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

3.4.1. Advisory list

This release includes the following advisories:

RHEA-2023:4577

Release of components for Red Hat OpenStack Platform 17.1 (Wallaby)

RHEA-2023:4578

Release of containers for Red Hat OpenStack Platform 17.1 (Wallaby)

RHEA-2023:4579

Red Hat OpenStack Platform 17.1 RHEL 9 deployment images

RHEA-2023:4580

Release of components for Red Hat OpenStack Platform 17.1 (Wallaby)

RHEA-2023:4581

Release of containers for Red Hat OpenStack Platform 17.1 (Wallaby)

RHSA-2023:4582

Moderate: Release of containers for Red Hat OpenStack Platform 17.1 director Operator

3.4.2. Bug fixes

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

BZ#1965308

Before this update, the Load-balancing service (octavia) could unplug a required subnet when you used different subnets from the same network as members' subnets. The members attached to this subnet were unreachable. With this update, the Load-balancing service does not unplug required subnets, and the load balancer can reach subnet members.

BZ#2007314

Before this update, instances with an emulated Trusted Platform Module (TPM) device could not be created due to an issue with the SElinux configuration in the **nova_libvirt** container. With this update, the deployment tooling configures SElinux correctly, which resolves the issue.

BZ#2066866

Even though the Panko monitoring service was deprecated, its endpoint still existed in the Identity service (keystone) after upgrading from RHOSP 16.2 to 17.1. With this update, the Panko service endpoint is cleaned up. However, Panko service users are not removed automatically. You must manually delete Panko service users with the command **openstack user delete panko**. There is no impact if you do not delete these users.

BZ#2073530

Support for the Windows Server 2022 guest operating system was not available in RHOSP 17.0 because it needs vTPM, and vTPM was not available due to an SElinux configuration issue. This issue has been fixed, and the Windows Server 2022 guest operating system is supported in RHOSP 17.1.

BZ#2080199

Before this update, services that were removed from the undercloud were not cleaned up during upgrades from RHOSP 16.2 to 17.0. The removed services remained in the OpenStack endpoint list even though they were not reachable or running. With this update, RHOSP upgrades include Ansible tasks to clean up the endpoints that are no longer required.

BZ#2089512

The multi-cell and multi-stack overcloud features were not available in RHOSP 17.0, due to a regression. The regressions have been fixed, and multi-cell and multi-stack deployments are supported in RHOSP 17.1.

BZ#2092444

Before this update, a bare-metal overcloud node was listed as active by the **metalsmith** tool even after being deleted. This happened in environments where the node naming scheme overlapped with the overcloud role naming scheme, which could result in the wrong node being unprovisioned during undeploy. Because the **metalsmith** tool uses the allocation name (hostname) first to lookup the status of bare-metal nodes, it was sometimes finding deleted nodes as still active.

With this update, nodes to be unprovisioned are now referenced by allocation name (hostname), which ensures that the correct node is always unprovisioned. The nodes are only referenced by node name if the hostname doesn't exist.

BZ#2097844

Before this update, the **overcloud config download** command failed with a traceback error because the command attempted to reach the Orchestration service (heat) to perform the download. The Orchestration service no longer persistently runs on the undercloud. With this update, the **overcloud config download** command is removed. Instead, you can use your **overcloud deploy** command with the **--stack-only** option.

BZ#2101846

Before this update, if secure RBAC was enabled, missing roles in the RHOSP deployment could cause Load-balancing service (octavia) API failures. In RHOSP 17.1 GA, this issue has been resolved.

BZ#2107580

Before this update, the shutdown script that director uses to stop **libvirtd** stored outdated **libvirt** container names from RHOSP versions before RHOSP 17.0, and instances did not shut down gracefully. With this update, the script stores correct **libvirt** container names, and instances are gracefully shut down when **libvirtd** is stopped.

BZ#2109616

Before this update, the Compute service was unable to determine the VGPU resource use because the mediated device name format changed in libvirt 7.7. With this update, the Compute service can now parse the new mediated device name format.

BZ#2116600

Before this update, the following libvirt internal error was sometimes raised during a successful live migration: "migration was active, but no RAM info was set". This caused the live migration to fail when it should have succeeded. With this update, when this libvirt internal error is raised, the live migration is signaled as complete in the libvirt driver and the live migration correctly succeeds.

BZ#2120145

Before this update, the low default value of the libvirt **max_client** parameter caused communication issues between libvirt and the Compute service (nova), which resulted in some failed operations, such as live migrations. With this update, you can customize the **max_client** parameter setting and increase its value to improve communication between libvirt and the Compute service.

BZ#2120767

The AMD SEV feature was not available in RHOSP 17.0, due to a known issue with the RHEL firmware definition file missing from some machine types. This issue has been fixed, and AMD SEV is supported in RHOSP 17.1.

BZ#2125610

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 17.1, this issue has been fixed and ICMP health monitors function correctly.

BZ#2125612

Before this update, users might have experienced the following warning message in the amphora log file of the Load-balancing service (octavia) when the load balancer was loaded with multiple concurrent sessions: **nf_conntrack: table full, dropping packet**. This error occurred if the amphora dropped Transport Control Protocol (TCP) flows and caused latency on user traffic. With this update, connection tracking (conntrack) is disabled for TCP flows in the Load-balancing service that uses amphora, and new TCP flows are not dropped. Conntrack is only required for User Datagram Protocol (UDP) flows.

BZ#2129207

Before this update, a network disruption or temporary unavailability of the Identity service (keystone) resulted in the nova-conductor service failing to start. With this update, the nova-conductor service logs a warning and continues startup in the presence of disruptions that are likely to be temporary. As a result, the nova-conductor service does not fail to start if transient issues like network disruptions or temporary unavailability of necessary services are encountered during startup.

BZ#2133027

The Alarming service (aodh) uses the deprecated gnocchi API to aggregate metrics, which results in incorrect metric measures of CPU usage in gnocchi. With this update, dynamic aggregation in gnocchi supports the ability to make re-aggregations of existing metrics and the ability to manipulate and transform metrics as required. CPU time in gnocchi is correctly calculated.

BZ#2133297

Before this update, the **openstack undercloud install** command launched the **openstack tripleo deploy** command, which created the **/home/stack/.tripleo/history** file with **root:root** as the owner. Subsequent deploy commands failed because of permission errors. With this update, the command creates the file with the **stack** user as the owner, and deploy commands succeed without permission errors.

BZ#2135548

Before this update, the **ironic-python-agent** did not correctly process the UEFI boot loader hint file, causing deployments to fail with RHEL 8.6 images in UEFI mode. With this update, you can now deploy RHEL 8.6 in UEFI mode.

BZ#2136302

This update allows node names longer than 62 bytes.

BZ#2140988

Before this update, a live migration might fail because the database did not update with the destination host details.
With this update, the instance host value in the database is set to the destination host during live migration.

BZ#2149216

Before this update, Open Virtual Network (OVN) load balancer health checks were not performed if you used Floating IPs (FIP) associated with the Load Balancer Virtual IP (VIP), and traffic was redirected to members in the Error state if the FIP was used.
With this update, if you use Floating IPs (FIP) associated with the Load Balancer Virtual IP (VIP), there is a new load balancer health check created for the FIP, and traffic is not redirected to members in the Error state.

BZ#2149221

Before this update, deployments with bonded interfaces did not complete because no value was set for the Ansible variable for OVS bonds, **bond_interface_ovs_options**. With this update, a default value has been set for the **bond_interface_ovs_options** Ansible variable.

BZ#2149339

Before this update, the cephadm-ansible logs in **/home/stack/config-download/overcloud/cephadm** were not rotated. The **cephadm_command.log** was appended for every overcloud deployment and increased in size. Also, for every **openstack overcloud ceph spec** operation, the log **/home/stack/ansible.log** was not rotated.
Now, dated logs are generated for every overcloud deployment, and every Ceph spec operation in the following format:

- **/home/stack/config-download/overcloud/cephadm/cephadm_command.log-<Timestamp>**.
- **/home/stack/ansible.log-<Timestamp>**.

BZ#2149468

Before this update, the Compute service (nova) processed a temporary error message from the Block Storage service (cinder) volume detach API, such as '504 Gateway Timeout', as an error. The Compute service failed the volume detach operation even though it succeeded but timed out on the Block Storage service side, leaving a stale block device mapping record in the Compute service database. With this update, the Compute service retries the volume detach call to the Block Storage service API if it receives an HTTP error that is likely to be temporary. Upon retry, if the volume attachment is no longer found, the Compute service processes the volume as already detached.

BZ#2149963

Before this update, the cephadm utility did not process child groups when building specification files from inventory. With this update, specification file generation processes child groups.

BZ#2151043

Before this update, the **openstack-cinder-volume-0** container, which is created by the Pacemaker bundle resource for the Block Storage service (cinder), mounted **/run** from the host. This mount path created the **.containerenv** file in the directory. When the **.containerenv** file exists,

subscription-manager fails because it evaluates that the command is executed inside a container. With this update, the mount path is updated so that Podman disables the creation of the **.containerenv** file, and **subscription-manager** executes successfully in a host that is running the **openstack-cinder-volume-0** container.

BZ#2152888

Before this update, the Service Telemetry Framework (STF) API health monitoring script failed because it depended on Podman log content, which was no longer available. With this update, the health monitoring script depends on the Podman socket instead of the Podman log, and API health monitoring operates normally.

BZ#2154343

Before this update, the disabling and enabling of network log objects in a security group was inconsistent. The logging of a connection was disabled as soon as one of the log objects in the security group associated with that connection was disabled. With this update, a connection is logged if any of the related enabled log objects in the security group allow it, even if one of those log objects becomes disabled.

BZ#2162632

Before this update, values of multi-value parameters were not populated correctly in the Alarming service (aodh) configuration because input to multi-value parameters was not considered as an array instead as a single value. With this update, you can set multiple values for a parameter and all values are populated in a configuration file.

BZ#2162756

Before this update, VLAN network traffic was centralized over the Controller nodes. With this update, if all the tenant provider networks that are connected to a router are of the VLAN/Flat type, that traffic is now distributed. The node that contains the instance sends the traffic directly.

BZ#2163815

Before this update, Open Virtual Network (OVN) load balancers on switches with **localnet** ports (Networking service [neutron] provider networks) did not work if traffic came from **localnet**. With this update, load balancers are not added to the logical switch associated with the provider network. This update forces Network Address Translation (NAT) to occur at the virtual router level instead of the logical switch level.

BZ#2164421

Before this update, the Compute service (nova) did not confidence-check the content of Virtual Machine Disk (VMDK) image files. By using a specially crafted VMDK image, it was possible to expose sensitive files on the host file system to guests booted with that VMDK image. With this update, the Compute service confidence checks VMDK files and forbids VMDK features that the leak behavior depends on. It is no longer possible to leak sensitive host file system contents using specially crafted VMDK files. This bug fix addresses CVE-2022-47951.

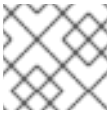

NOTE

Red Hat does not support the VMDK image file format in RHOSP.

BZ#2164677

Before this update, the iptables rule for the heat-cfn service contained the incorrect TCP port number. Users could not access the heat-cfn service endpoint if SSL was enabled for public endpoints. With this update, the TCP port number is correct in the iptables rule. Users can access the heat-cfn service endpoint, even if SSL is enabled for public endpoints.

BZ#2167161

Before this update, the default value of **rgw_max_attr_size** was 256, which created issues for OpenShift on OpenStack when uploading large images. With this update, the default value of **rgw_max_attr_size** is 1024.
You can change the value by adding the following configuration to an environment file that you include in your overcloud deployment:

```
parameters_default:
  CephConfigOverrides:
    rgw_max_attr_size: <new value>
```

BZ#2167431

Before this update, the collectd hugepages plugin would report a failure message when attempting to access a new file in Red Hat Enterprise Linux (RHEL) 9 called **demote**. Now, collectd avoids reading this file and the failure message is suppressed.

BZ#2169303

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. The IPMI agent container is executed with the **--privilege** flag to execute **ipmitool** commands on the host. The data collection service (ceilometer) can now capture power metrics.

BZ#2169349

Before this update, instances lost communication with the ovn-metadata-port because the load balancer health monitor replied to the ARP requests for the OVN metadata agent's IP, causing the request going to the metadata agent to be sent to another MAC address. With this update, the ovn-controller conducts back-end checks by using a dedicated port instead of the ovn-metadata-port. When establishing a health monitor for a load balancer pool, ensure that there is an available IP in the VIP load balancer's subnet. This port is distinct for each subnet, and various health monitors in the same subnet can reuse the port. Health monitor checks no longer impact ovn-metadata-port communications for instances.

BZ#2172063

Before this update, the **openstack overcloud ceph deploy** command could fail during the **apply spec** operation if the chrony NTP service was down. With this update, the chrony NTP service is enabled before the **apply spec** operation.

BZ#2172582

Before this update, the **create pool** operation failed because the podman command used **/etc/ceph** as the volume argument. This argument does not work for Red Hat Ceph Storage version 6 containers. With this update, the podman command uses **/var/lib/ceph/$FSID/config/** as the first volume argument and **create pool** operations are successful.

BZ#2173101

Before this update, when users deployed Red Hat Ceph Storage in a tripleo-ipa context, a **stray hosts** warning showed in the cluster for the Ceph Object Gateway (RADOS Gateway [RGW]). With this update, during a Ceph Storage deployment, you can pass the option **--tld** in a tripleo-ipa context to use the correct hosts when you create the cluster.

BZ#2173575

Before this update, a flooding issue occurred when an instance, associated with a provider network with disabled port security, attempted to reach IPs on the provider network that were not recognized by OpenStack. This flooding occurred because the forwarding database (FDB) table was not learning MAC addresses. This update uses a new option in OVN to enable the learning of IPs in the FDB table. There is currently no aging mechanism for the FDB table. But you can clean up the FDB table periodically, to prevent the occurrence of scaling issues caused by the size of this table.

BZ#2174632

Before this update, a regression in the network configuration for OVS interfaces negatively impacted network performance. With this update, the **os-vif** OVS plugin has been enhanced to improve network performance on the OVS interfaces of non-Windows instances.

IMPORTANT

This update takes effect when the instance interface is recreated. If you change this value for an existing port, you must hard reboot the instance or perform a live migration for the update to take effect.

BZ#2178618

Before this update, a security group logging enhancement introduced an issue where log objects could not be deleted at the same time as security groups. This action caused an internal server error. With this update, the **db_set** function that modifies the northbound database entries does not fail if the row that is requested does not exist any more.

BZ#2179071

Before this update, the collectd plugin libpodstats could not gather metrics because the Cgroup path to Ceph containers changed in RHEL 9 from **/sys/fs/cgroup/machine.slice** to **/sys/fs/cgroup/system.slice/system-ceph<FSID>**. With this update, libpodstats can now parse CPU and memory metrics from cgroups under the new path.

BZ#2180933

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/**. Because of this issue, the rsyslog service could not locate Pacemaker log files. With this update, the Pacemaker log path is changed from **/var/log/pacemaker/** to **/var/log/host/pacemaker/**.

BZ#2181107

Before this update the **NetworkDeploymentAction** parameter was internally overridden and the deployment process would always configures the network interfaces. As a result, the network interfaces were always configured during deployment regardless of the value of the **NetworkDeploymentAction** parameter. With this update the **NetworkDeploymentAction** parameter works as expected, and by default the configuration of networking interfaces is skipped for nodes that are already deployed.

BZ#2185163

Before this update, existing puppet containers were reused during deployment. The deployment process did not check the return code from the puppet commands executed within the container, which meant that any puppet task failures were ignored during deployment. This resulted in reporting a successful deployment even when some puppet execution tasks failed. With this update, puppet containers are recreated for every deployment. If a puppet execution task fails, the deployment stops and reports the failure.

BZ#2188252

Before this update, the 'openstack tripleo container image prepare' command failed because there were incorrect Ceph container tags in the **container_image_prepare_defaults.yaml** file. With this update, the correct Ceph container tags are in the YAML file, and the 'openstack tripleo container image prepare' command is successful.

BZ#2196288

Before this update, if you upgraded your operating system from RHEL 7.x to RHEL 8.x, or from RHEL 8.x to RHEL 9.x, and ran a Leapp upgrade with the **--debug** option, the system remained in the **early console in setup code** state and did not reboot automatically. With this update, the

UpgradeLeappDebug parameter is set to **false** by default. Do not change this value in your templates.

BZ#2203238

Before this update, for the nova-compute log to record os-brick privileged commands for debugging purposes, you had to apply the workaround outlined in https://access.redhat.com/articles/5906971. This update makes the workaround redundant and provides a better solution that separates logging by the nova-compute service so that the privileged commands of os-brick are logged at the debug level but the privileged commands of nova are not.

BZ#2207991

Before this update, secure role-based access control (SRBAC) and the **NovaShowHostStatus** parameter used the same policy key titles. If you configured both SRBAC and **NovaShowHostStatus**, the deployment failed with a conflict. With this update, the policy key for **NovaShowHostStatus** is changed and there are no related conflicts in deployments.

BZ#2210062

Before this update, in RHOSP 17.1 environments that use RHOSP dynamic routing, there was a known issue where the default value of the Autonomous System Number (ASN) used by the OVN BGP agent differed from the ASN used by FRRouting (FRR).
In 17.1 GA, this issue is resolved. The **FrrOvnBgpAgentAsn** and **FrrBgpAsn** default values are valid and can be used without needing to modify them.

BZ#2211691

Before this update, the Bare Metal Provisioning service (ironic) was unable to detach a Block Storage service (cinder) volume from a physical bare metal node. This volume detachment is required to tear down physical machines that have an instance deployed on them by using the boot from volume functionality. With this update, the Bare Metal Provisioning service (ironic) can detach a volume from a physical bare metal node to automatically tear down these physical machines.

BZ#2211849

Before this update, a bug in the library **pyroute2** caused environments that used RHOSP dynamic routing to fail to advertise new routes and to lose connectivity with new or migrated instances, new load balancers, and so on. In RHOSP 17.1 GA, a newer version of **pyroute2** resolves this issue.

BZ#2214259

Before this update, in an environment that had been migrated from the OVS mechanism driver to the OVN mechanism driver, an instance with a trunk port could become inaccessible after an operation such as a live migration. Now, you can live migrate, shutdown, or reboot instances with a trunk port without issues after migration to the OVN mechanism driver.

BZ#2215936

Before this update, creating an instance with virtual functions (VF) could fail in an environment that had been migrated from ML2/OVS with SR-IOV to ML2/OVN. You can now create instances with VFs after migration.

BZ#2216130

Currently, **puppet-ceilometer** does not populate the **tenant_name_discovery** parameter in the data collection service (ceilometer) configuration on Compute nodes. This causes the **Project name** and **User name** fields to not be identified. Currently, there is no workaround for this issue.

BZ#2219765

Before this update, the **pam_loginuid** module was enabled in some containers. This prevented crond from executing some tasks, such as **db purge,** inside of those containers. Now, **pam_loginuid** is removed and the containerized **crond** process runs all periodic tasks.

3.4.3. Enhancements

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

BZ#1369007

Cloud users can launch instances that are protected with UEFI Secure Boot when the overcloud contains UEFI Secure Boot Compute nodes. For information on creating an image for UEFI Secure Boot, see Creating an image for UEFI Secure Boot . For information on creating a flavor for UEFI Secure Boot, see "UEFI Secure Boot" in Flavor metadata.

BZ#1581414

Before this release, **NovaHWMachineType** could not be changed for the lifetime of a RHOSP deployment because the machine type of instances without a **hw_machine_type** image property would use the newly configured machine types after a hard reboot or migration. Changing the underlying machine type for an instance could break the internal ABI of the instance.
With this release, when launching an instance the Compute service records the instance machine type within the system metadata of the instance. Therefore, it is now possible to change the **NovaHWMachineType** during the lifetime of a RHOSP deployment without affecting the machine type of existing instances.

BZ#1619266

This update introduces the security group logging feature. To monitor traffic flows and attempts into and out of an instance, you can configure the 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 instance, 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.

You can use security group logging for both stateful and stateless security groups.

Logged events are stored on the Compute nodes that host the instances, in the file **/var/log/containers/stdouts/ovn_controller.log**.

BZ#1666804

With this update, the **cinder-backup** service can now be deployed in Active/Active mode.

BZ#1672972

This enhancement helps cloud users determine if the reason they are unable to access an "ACTIVE" instance is because the Compute node that hosts the instance is unreachable. RHOSP administrators can now configure the following parameters to enable a custom policy that provides a status in the **host_status** field to cloud users when they run the **openstack show server details** command, if the host Compute node is unreachable:

- **NovaApiHostStatusPolicy**: Specifies the role the custom policy applies to.
- **NovaShowHostStatus**: Specifies the level of host status to show to the cloud user, for example, "UNKNOWN".

BZ#1693377

With this update, an instance can have a mix of shared (floating) CPUs and dedicated (pinned) CPUs instead of only one CPU type. RHOSP administrators can use the **hw:cpu_policy=mixed** and **hw_cpu_dedicated_mask** flavor extra specs to create a flavor for instances that require a mix of shared CPUs and dedicated CPUs.

BZ#1701281

In RHOSP 17.1, support is available for cold migrating and resizing instances that have vGPUs.

BZ#1720404

With this update, you can configure your RHOSP deployment to count the quota usage of cores and RAM by querying placement for resource usage and instances from instance mappings in the API database, instead of counting resources from separate cell databases. This makes quota usage counting resilient to temporary cell outages or poor cell performance in a multi-cell environment. Set the following configuration option to count quota usage from placement:

```
parameter_defaults:
  ControllerExtraConfig:
    nova::config::nova_config:
      quota/count_usage_from_placement:
        value: 'True'
```

BZ#1761861

With this update, you can configure each physical GPU on a Compute node to support a different virtual GPU type.

BZ#1761903

On RHOSP deployments that use a routed provider network, you can now configure the Compute scheduler to filter Compute nodes that have affinity with routed network segments, and verify the network in placement before scheduling an instance on a Compute node. You can enable this feature by using the **NovaSchedulerQueryPlacementForRoutedNetworkAggregates** parameter.

BZ#1772124

With this update, you can use the new **NovaMaxDiskDevicesToAttach** heat parameter to specify the maximum number of disk devices that can be attached to a single instance. The default is unlimited (-1). For more information, see Configuring the maximum number of storage devices to attach to one instance.

BZ#1782128

In RHOSP 17.1, a RHOSP administrator can provide cloud users the ability to create instances that have emulated virtual Trusted Platform Module (vTPM) devices. RHOSP only supports TPM version **2.0**.

BZ#1793700

In RHOSP 17.1, a RHOSP administrator can declare which custom physical features and consumable resources are available on the RHOSP overcloud nodes by modeling custom traits and inventories in a YAML file, **provider.yaml**.

BZ#1827598

This RHOSP release introduces support of the OpenStack stateless security groups API.

BZ#1857652

With this update, deployments of RHOSP with trunk ports are fully supported for migration from ML2/OVS to ML2/OVN.

BZ#1873409

On RHOSP deployments that are configured for OVS hardware offload and to use ML2/OVN, and that have Compute nodes with VirtIO data path acceleration (VDPA) devices and drivers and Mellanox NICs, you can enable VDPA support for enterprise workloads. When VDPA support is enabled, your cloud users can create instances that use VDPA ports. For more information, see Configuring VDPA Compute nodes to enable instances that use VDPA ports and Creating an instance with a VDPA interface.

BZ#1873707

With this update, you can use the validation framework in the workflow of backup and restore procedures to verify the status of the restored system. The following validations are included:

- **undercloud-service-status**
- **neutron-sanity-check**
- **healthcheck-service-status**
- **nova-status**
- **ceph-health**
- **check-cpu**
- **service-status**
- **image-serve**
- **pacemaker-status**
- **validate-selinux**
- **container-status**

BZ#1883554

With this update, a RHOSP administrator can create a flavor that has a **socket** PCI NUMA affinity policy. You can use this policy to create an instance that requests a PCI device only when at least one of the instance NUMA nodes has affinity with a NUMA node in the same host socket as the PCI device.

BZ#1888788

With this update, the Shared File Systems service (manila) API supports a project-scoped 'reader' role. Users with the 'reader' role can send GET requests to the service, but they cannot make any other kind of request. You can enable this feature by using the **environments/enable-secure-rbac.yaml** environment file included with director. You can use the 'reader' role to create audit users for humans and automation and to perform read-only interactions safely with OpenStack APIs.

BZ#1898349

With this update, the Block Storage (cinder) backup service supports the zstd data compression algorithm.

BZ#1903914

With this update, the Block Storage (cinder) backup service supports the S3 back end.

BZ#1947377

With this update, the RHOSP Orchestration service (heat) dashboard shows template default values. Previously, the heat dashboard had the default values hidden, which was sometimes confusing for users. This update ensures that those default values are visible to the user in the heat dashboard and removes any confusion that was caused when they were hidden.

BZ#1962500

With this update, you can configure the collectd logging source in TripleO Heat Templates. The default value matches the default logging path.

BZ#1986025

With this update, Block Storage service (cinder) supports NVMe over TCP (NVMe/TCP) drivers, for Compute nodes that are running RHEL 9.

BZ#2005495

This enhancement allows cloud administrators to specify an Availability Zone (AZ) by storage back end through director when configuring the Shared File Systems service (manila) back-end storage. With this update, administrators can use an AZ annotation to logically separate storage provisioning requests and to denote failure domains. AZs configured by administrators are exposed by the Shared File Systems service to end users. End users can request that their workloads be scheduled to specific AZs based on their needs. When configuring multiple storage back ends, administrators might want to tag each back end to different AZs as opposed to denoting a single AZ for all back ends.

Director has new options to denote the storage AZs. Each option corresponds to a supported storage back-end driver. For more information about AZs, see *Configuring persistent storage*.

BZ#2008969

With this update, cloud administrators can bring shares that are created outside the Shared File Systems service (manila) under the management of the Shared file Systems service. Cloud administrators can also remove shares from the Shared File Systems service without deleting them. Note that the CephFS driver does not support this feature. You can use this manage/unmanage functionality when commissioning, decommissioning, or migrating storage systems, or to take shares offline temporarily for maintenance.

BZ#2016660

Upgrades from Red Hat OpenStack Platform (RHOSP) 16.2 to RHOSP 17.1 are supported. The RHOSP upgrade and the operating system upgrade are now separated into two distinct phases. You upgrade RHOSP first, then you upgrade the operating system.

BZ#2026385

With this update, you can configure **fence_watchdog** that uses **sbd**, like other fencing devices via tripleo, by defining the respective fencing resource:

```
parameter_defaults:
  EnableFencing: true
  FencingConfig:
    devices:
    - agent: fence_watchdog
      host_mac: "52:54:00:74:f7:51"
```

As an operator, you must enable **sbd** and set the watchdog timeout:

```
parameter_defaults:
  ExtraConfig:
    pacemaker::corosync::enable_sbd: true
    tripleo::fencing::watchdog_timeout: 20
```

BZ#2033811

The Shared File System service (manila) now supports using Pure Storage Flashblade system as a back end. Refer to the Red Hat ecosystem catalog to find the vendor's certification and installation documentation.

BZ#2060758

In Red Hat OpenStack Platform (RHOSP) 17.1, the RHOSP Load-balancing service (octavia) supports the rsyslog over TCP protocol for Amphora log offloading. With this enhancement you can

redirect log messages to a secondary rsyslog server if the primary server becomes unavailable. For more information, see Chapter 5. Managing Load-balancing service instance logs in the *Configuring load balancing as a service* guide.

BZ#2066349

With this enhancement, the LVM volumes installed by the **overcloud-hardened-uefi-full.qcow2** whole disk overcloud image are now backed by a thin pool. The volumes are still grown to consume the available physical storage, but are not over-provisioned by default.
The benefits of thin-provisioned logical volumes:

- If a volume fills to capacity, the options for manual intervention now include growing the volume to over-provision the physical storage capacity.
- The RHOSP upgrades process can now create ephemeral backup volumes in thin-provisioned environments.

BZ#2069624

The Red Hat OpenStack Platform (RHOSP) snapshot and revert feature is based on the Logical Volume Manager (LVM) snapshot functionality and is intended to revert an unsuccessful upgrade or update. Snapshots preserve the original disk state of your RHOSP cluster before performing an upgrade or an update. You can then remove or revert the snapshots depending on the results. If an upgrade completed successfully and you do not need the snapshots anymore, remove them from your nodes. If an upgrade fails, you can revert the snapshots, assess any errors, and start the upgrade procedure again. A revert leaves the disks of all the nodes exactly as they were when the snapshot was taken.

BZ#2074896

Previously, the Open vSwitch (OVS) bond **balance-tcp** mode was only available in RHOSP as a technology preview. Because of L4 hashing re-circulation issues, the mode was not recommended for production. The issues have been resolved and you can use the OVS bond **balance-tcp** mode. You must set **lb-output-action=true** to use **balance-tcp** mode.

BZ#2086688

RHOSP 17.1 GA supports the offloading of OpenFlow flows to hardware with the connection tracking (conntrack) module. For more information, see Configuring components of OVS hardware offload in *Configuring network functions virtualization*.

BZ#2097931

In RHOSP 17.1, you can live migrate, unshelve and evacuate an instance that uses a port that has resource requests, such as a guaranteed minimum bandwidth QoS policy.

BZ#2104522

With this update, live migration now uses multichassis Open Virtual Network (OVN) ports to optimize the migration procedure and significantly reduce network downtime for VMs during migration in particular scenarios.

BZ#2106406

This update introduces the script **neutron-remove-duplicated-port-bindings** to fix an issue that sometimes affected the handling of failed live migrations.
If a live migration fails, the Compute service (Nova) reverts the migration. The migration reversal implies deleting any object created in the database or in the destination compute node.

However, in some cases after the reversal of a failed live migration, ports were left with duplicate port bindings.

The **neutron-remove-duplicated-port-bindings** script finds duplicate port bindings and deletes the inactive bindings. You can run the script if a failed live migration results in duplicate port bindings.

BZ#2111528

With this update, the default Ceph container image is based on Red Hat Ceph Storage 6 instead of Red Hat Ceph Storage 5.

BZ#2122209

This update adds the **validation file** command to the Validation Framework CLI. This command allows you to supply a file with validations by name, group, category and product for a validation run. Now, you can run 'validation file <path_to_file>', and keep the chosen validations for reruns at a later time.

BZ#2124309

With this enhancement, operators can enable the run_arping feature for Pacemaker-managed virtual IPs (VIPs), so that the cluster preemptively checks for duplicate IPs.
To do this, you must add the following configuration to the environment file:

```
ExtraConfig:
  pacemaker::resource::ip::run_arping: true
```

If a duplicate is found, the following error is logged in the **/var/log/pacemaker/pacemaker.log** file:

```
Sep 07 05:54:54  IPaddr2(ip-172.17.3.115)[209771]:    ERROR: IPv4 address collision
172.17.3.115 [DAD]
Sep 07 05:54:54  IPaddr2(ip-172.17.3.115)[209771]:    ERROR: Failed to add 172.17.3.115
```

BZ#2138238

With this update, you deploy two separate instances of the Image service (glance) API. The instance that is accessible to OpenStack tenants is configured to hide image location details, such as the direct URL of an image or whether the image is available in multiple locations. The second instance is accessible to OpenStack administrators and OpenStack services, such as the Block Storage service (cinder) and the Compute service (nova). This instance is configured to provide image location details. This enhancement addresses the recommendations of OSSN-0090 and CVE-2022-4134. With this update, a malicious user cannot leverage the location details of an image to upload an altered image.

BZ#2152877

This enhancement adds OVN security group logging to the Networking service (neutron) for the reply packets of a network connection. The ovn-controller log files now log the full network connection.

BZ#2165501

Starting with Red Hat OpenStack Platform (RHOSP) 17.1, in ML2/OVN deployments, you can enable minimum bandwidth and bandwidth limit egress policies for hardware offloaded ports. You cannot enable ingress policies for hardware offloaded ports. For more information, see Configuring the Networking service for QoS policies.

BZ#2187255

With this update, you can add project and user name fields to outgoing data collection service (ceilometer) metrics. Previously, cloud administrators had to rely on UUIDs of projects and users to identify tenants. Now you can view a list of projects and user names, not UUIDs.

NOTE

This feature is not available to use with gnocchi or Service Telemetry Framework (STF).

3.4.4. Technology previews

The items listed in this section are provided as Technology Previews in this release of 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#1813561

With this update, the Load-balancing service (octavia) supports HTTP/2 load balancing by using the Application Layer Protocol Negotiation (ALPN) for listeners and pools that are enabled with Transport Layer Security (TLS). The HTTP/2 protocol improves performance by loading pages faster.

BZ#1848407

In RHOSP 17.1, a technology preview is available for the Stream Control Transmission Protocol (SCTP) in the Load-balancing service (octavia). Users can create SCTP listeners and attach SCTP pools in a load balancer.

BZ#2057921

In RHOSP 17.1, a technology preview is available for creating load balancers over an IPv6 management network. Using a private IPv6 management network for the Load-balancing service (octavia) may simplify edge deployments.

BZ#2217663

In RHOSP 17.1, a technology preview is available for the VF-LAG transmit hash policy offload that enables load balancing at NIC hardware for offloaded traffic/flows. This hash policy is only available for layer3+4 base hashing.

3.4.5. Release notes

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

BZ#2072644

This enhancement allows users to upgrade from RHOSP 16.2 to RHOSP 17.1 and keep the Red Hat Enterprise Linux (RHEL) 8 based operating systems on the Compute nodes, in combination with nodes running RHEL 9.
Control plane nodes and Storage nodes must be upgraded. The default behavior is that all nodes are upgraded to RHEL 9 unless explicitly configured otherwise.

BZ#2081641

If you are using a Red Hat OpenStack Platform (RHOSP) environment that is running RHOSP 16.2.4 or later, you can upgrade directly to RHOSP 17.1.

BZ#2224523

In RHOSP networking environments, when creating a VM instance, do not bind the instance to a virtual port (vport). Instead, use a port whose IP address is not a member of another port's allowed address pair.
Binding a vport to an instance prevents the instance from spawning and produces an error message similar to the following:

```
WARNING nova.virt.libvirt.driver [req-XXXX - - - default default] [instance: XXXXXXXXX] Timeout waiting for [('network-vif-plugged', 'XXXXXXXXXX')] for instance with vm_state building and task_state spawning.: eventlet.timeout.Timeout: 300 seconds
```

3.4.6. Known issues

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

BZ#2108212

If you use IPv6 to connect to instances during migration to the OVN mechanism driver, connection to the instances might be disrupted for up to several minutes when the ML2/OVS services are stopped. To avoid this, use IPv4 instead.
The router advertisement daemon **radvd** for IPv6 is stopped during migration to the OVN mechanism driver. While **radvd** is stopped, router advertisements are no longer broadcast. This broadcast interruption results in instance connection loss over IPv6. IPv6 communication is automatically restored once the new ML2/OVN services start.

To avoid the potential disruption, use IPv4 instead.

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 TX flow in HW. CX-5 cannot pop VLAN on TX path, which prevents offloading the flow with pop_vlan to the HW.

BZ#2109985

Currently, in ML2/OVS deployments, Open vSwitch (OVS) does not support offloading OpenFlow rules that have the **skb_priority**, **skb_mark**, or output queue fields set. These fields are required for Quality of Service (QoS) support for virtio ports.
If you set a minimum bandwidth rule for a virtio port, the Networking service (neutron) OVS agent marks the traffic of this port with a Packet Mark field. 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 that no traffic can be offloaded.

Workaround: If your environment includes OVS hardware offload ports, disable packet marking in the nodes that require hardware offloading. When you disable packet marking, it is not possible to set rate limiting rules for virtio ports. However, differentiated services code point (DSCP) marking rules are still available.

In the configuration file, set the **disable_packet_marking** flag to **true**. When 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#2126725

Hard-coded certificate location operates independently of user-provided values. During deployment with custom certificate locations, services do not retrieve information from API endpoints because Transport Layer Security (TLS) verification fails.

BZ#2143874

In RHOSP 17.1, when the DNS service (designate) is deployed, Networking service (neutron) ports created on the undercloud are not deleted when the overcloud is deleted. These ports do not cause operational problems when the overcloud is recreated with or without the DNS service.
Workaround: After the overcloud has been deleted, manually remove the ports by using the **openstack port delete** command.

BZ#2144492

If you migrate a RHOSP 17.1.0 ML2/OVS deployment with distributed virtual routing (DVR) to ML2/OVN, the floating IP (FIP) downtime that occurs during ML2/OVN migration can exceed 60 seconds.

BZ#2160481

In RHOSP 17.1 environments that use BGP dynamic routing, there is currently a known issue where floating IP (FIP) port forwarding fails.
When FIP port forwarding is configured, packets sent to a specific destination port with a destination IP that equals the FIP are redirected to an internal IP from a RHOSP Networking service (neutron) port. This occurs regardless of the protocol that is used: TCP, UDP, and so on.

When BGP dynamic routing is configured, the routes to the FIPs used to perform FIP port forwarding are not exposed, and these packets cannot reach their final destinations.

Currently, there is no workaround.

BZ#2163477

In RHOSP 17.1 environments that use BGP dynamic routing, there is currently a known issue affecting instances connected to provider networks. The RHOSP Compute service cannot route packets sent from one of these instances to a multicast IP address destination. Therefore, instances subscribed to a multicast group fail to receive the packets sent to them. The cause is that BGP multicast routing is not properly configured on the overcloud nodes. Currently, there is no workaround.

BZ#2167428

During a new deployment, the keystone service is often not available when the agent-notification service is initializing. This prevents ceilometer from discovering the gnocchi endpoint. As a result, metrics are not sent to gnocchi.

BZ#2178500

If a volume refresh fails when using the nova-manage CLI, this causes the instance to stay in a locked state.

BZ#2180542

The Pacemaker-controlled **ceph-nfs** resource requires a runtime directory to store some process data. The directory is created when you install or upgrade RHOSP. Currently, a reboot of the Controller nodes removes the directory, and the **ceph-nfs** service does not recover when the Controller nodes are rebooted. If all Controller nodes are rebooted, the **ceph-nfs** service fails permanently.
Workaround: If you reboot a Controller node, log into the Controller node and create a **/var/run/ceph** directory:

$ mkdir -p /var/run/ceph

Repeat this step on all Controller nodes that have been rebooted. If the **ceph-nfs-pacemaker** service has been marked as failed, after creating the directory, execute the following command from any of the Controller nodes:

$ pcs resource cleanup

BZ#2180883

Currently, Logrotate archives all log files once a day and Rsyslog stops sending logs to Elasticsearch
Workaround: Add "RsyslogReopenOnTruncate: true" to your environment file during deployment so that Rsyslog reopens all log files on log rotation.
Currently, RHOSP 17.1 uses an older puppet-rsyslog module with an incorrectly configured Rsyslog.
Workaround: Manually apply patch [1] in **/usr/share/openstack-tripleo-heat-**

templates/deployment/logging/rsyslog-container-puppet.yaml before deployment to configure Rsyslog correctly.

BZ#2182371

There is currently a known issue with guest instances that use Mellanox ConnectX-5, ConnectX-6, and Bluefield-2 NICs with offload (switchdev) ports. It takes a long time to initialize the system when you reboot the operating system from the guest directly, for example, by using the command **sudo systemctl reboot --reboot-arg=now**. If the instance is configured with two Virtual Functions (VFs) from the same Physical Function (PF), the initialization of one of the VFs might fail and cause a longer initialization time.
Workaround: Reboot the guest instance in a timely manner by using the OpenStack API instead of rebooting the guest instance directly.

BZ#2183793

Overcloud node provisioning may fail for NFV deployments on some AMD platforms in UEFI boot mode on RHOSP 17.1, when using the following BIOS configuration:

- Boot Mode: UEFI
- Hard-disk Drive Placeholder: Enabled
  Workaround: Set **Hard-disk Drive Placeholder** to **Disabled**. For information on how to assess each BIOS attribute for your NFV deployment on AMD platforms in UEFI boot mode, see the reference guide for your hardware.

BZ#2184834

The Block Storage API supports the creation of a Block Storage multi-attach volume by passing a parameter in the volume-create request, even though this method of creating multi-attach volume has been deprecated for removal because it is unsafe and can lead to data loss when creating a multi-attach volume on a back end that does not support multi-attach volumes. Workaround: create a multi-attach volume by using a multi-attach volume-type, which is the only method of creating multi-attach volumes provided by the **openstack** and **cinder** CLI.

BZ#2185897

In ML2/OVN deployments, do not use live migration on instances that use trunk ports. On instances that use trunk ports, live migration can fail due to the flapping of the instance's subport between the Compute nodes. For instances that have trunk ports, use cold migration instead.

BZ#2192913

In RHOSP environments with ML2/OVN or ML2/OVS that have DVR enabled and use VLAN tenant networks, east/west traffic between instances connected to different tenant networks is flooded to the fabric.
As a result, packets between those instances reach not only the Compute nodes where those instances run, but also any other overcloud node.

This could cause an impact on the network and it could be a security risk because the fabric sends traffic everywhere.

This bug will be fixed in a later FDP release. You do not need to perform a RHOSP update to obtain the FDP fix.

BZ#2193388

The Dashboard service (horizon) is currently configured to validate client TLS certificates by default, which breaks the Dashboard service on all TLS everywhere (TLS-e) deployments.
Workaround:

1. Add the following configuration to an environment file:

```
parameter_defaults:
  ControllerExtraConfig:
    horizon::ssl_verify_client: none
```

2. Add the environment file to the stack with your other environment files and deploy the overcloud:

```
(undercloud)$ openstack overcloud deploy --templates \
 -e [your environment files] \
 -e /home/stack/templates/<environment_file>.yaml
```

BZ#2196291

Currently, custom SRBAC rules do not permit list policy rules to non-admin users. As a consequence, non-admin users can not list or manage these rules. Current workarounds include either disabling SRBAC, or modifying the SRBAC custom rule to permit this action.

BZ#2203785

Currently, there is a permission issue that causes collectd sensubility to stop working after you reboot a baremetal node. As a consequence, sensubility stops reporting container health. Workaround: After rebooting an overcloud node, manually run the following command on the node: **sudo podman exec -it collectd setfacl -R -m u:collectd:rwx /run/podman**

BZ#2203857

A known issue in the Ceph RADOS Gateway component in Red Hat Ceph Storage (RHCS) 6.0 causes authorization with Identity service (keystone) tokens to fail. This issue is not manifest in RHCS 6.1, which is supported in RHOSP 17.1.

BZ#2210030

There is currently a known issue where custom SRBAC rules do not permit list shared security groups to non-administrative users that are not rule owners. This causes shared security groups and rules to not be managed properly by non-administrative users that are not rule owners. Workaround: Disable custom SRBAC rules or modify the custom rules to permit any user to manage the rules.

BZ#2210319

Currently, the Retbleed vulnerability mitigation in RHEL 9.2 can cause a performance drop for Open vSwitch with Data Plane Development Kit (OVS-DPDK) on Intel Skylake CPUs.
This performance regression happens only if C-states are disabled in the BIOS, hyper-threading is enabled, and OVS-DPDK is using only one hyper-thread of a given core.

Workaround: Assign both hyper-threads of a core to OVS-DPDK or to SRIOV guests that have DPDK running as recommended in the NFV configuration guide.

BZ#2213126

The logging queue that buffers excess security group log entries sometimes stops accepting entries before the specified limit is reached. As a workaround, you can set the queue length higher than the number of entries you want it to hold.
You can set the maximum number of log entries per second with the parameter **NeutronOVNLoggingRateLimit**. If the log entry creation exceeds that rate, the excess is buffered in a queue up to the number of log entries that you specify in **NeutronOVNLoggingBurstLimit**.

The issue is especially evident in the first second of a burst. In longer bursts, such as 60 seconds, the rate limit is more influential and compensates for burst limit inaccuracy. Thus, the issue has the greatest proportional effect in short bursts.

Workaround: Set **NeutronOVNLoggingBurstLimit** at a higher value than the target value. Observe and adjust as needed.

BZ#2215053

In RHOSP 17.1 environments that use Border Gateway Protocol (BGP) dynamic routing, there is currently a known issue where the FRRouting (FRR) container fails to deploy. This failure occurs because the RHOSP director deploys the FRR container before the container image prepare task finishes. Workaround: In your heat templates, ensure that the **ContainerImagePrepare** precedes the **overcloud deploy** command.

BZ#2216021

RHOSP 17.1 with the OVN mechanism driver does not support logging of flow events per port or the use of the **--target** option of the **network log create** command.
RHOSP 17.1 supports logging of flow events per security groups, using the **--resource** option of the **network log create** command. See "Logging security group actions" in *Configuring Red Hat OpenStack Platform networking*.

BZ#2217867

There is currently a known issue on Nvidia ConnectX-5 and ConnectX-6 NICs, when using hardware offload, where some offloaded flows on a PF can cause transient performance issues on the associated VFs. This issue is specifically observed with LLDP and VRRP traffic.

BZ#2219574

The data collection service (ceilometer) does not provide a default caching back end, which can cause some services to be overloaded when polling for metrics.

BZ#2219603

In RHOSP 17.1 GA, the DNS service (designate) is misconfigured when secure role-based access control (sRBAC) is enabled. The current sRBAC policies contain incorrect rules for designate and must be corrected for designate to function correctly.
Workaround: Apply the following patch on the undercloud server and redeploy the overcloud:

https://review.opendev.org/c/openstack/tripleo-heat-templates/+/888159

BZ#2219830

In RHOSP 17.1, there is a known issue of transient packet loss where hardware interrupt requests (IRQs) are causing non-voluntary context switches on OVS-DPDK PMD threads or in guests running DPDK applications.
This issue is the result of provisioning large numbers of VFs during deployment. VFs need IRQs, each of which must be bound to a physical CPU. When there are not enough housekeeping CPUs to handle the capacity of IRQs, **irqbalance** fails to bind all of them and the IRQs overspill on isolated CPUs.

Workaround: You can try one or more of these actions:

- Reduce the number of provisioned VFs to avoid unused VFs remaining bound to their default Linux driver.
- Increase the number of housekeeping CPUs to handle all IRQs.
- Force unused VF network interfaces down to avoid IRQs from interrupting isolated CPUs.
- Disable multicast and broadcast traffic on unused, down VF network interfaces to avoid IRQs from interrupting isolated CPUs.

BZ#2220808

In RHOSP 17.1, there is a known issue where the data collection service (ceilometer) does not report airflow metrics. This problem is caused because the data collection service is missing a gnocchi resource type, **hardware.ipmi.fan**. Currently, there is no workaround.

BZ#2220887

The data collection service (ceilometer) does not filter separate power and current metrics.

BZ#2222543

Currently, when a bootstrap Controller node is replaced, the OVN database cluster is partitioned: with two database clusters for both the northbound and southbound databases. This situation makes instances unusable.
To find the name of the bootstrap Controller node, run the following command:

```
ssh tripleo-admin@CONTROLLER_IP "sudo hiera -c /etc/puppet/hiera.yaml
pacemaker_short_bootstrap_node_name"
```

Workaround: Perform the steps described in Red Hat KCS solution 7024434: Recover from partitioned clustered OVN database.

BZ#2222589

There is currently a known issue with the upgrade from RHOSP 16.2 to 17.1, where the director upgrade script stops executing when upgrading Red Hat Ceph Storage 4 to 5 in a director-deployed Ceph Storage environment that uses IPv6. Workaround: Apply the workaround from Red Hat KCS solution 7027594: Director upgrade script stops during RHOSP upgrade when upgrading RHCS in director-deployed environment that uses IPv6

BZ#2222605

In RHOSP 17.1, there is a known issue for security group log entries. When events occur in short time intervals of each other, the related security group log entries can be listed in an incorrect order. This is caused by how the OVN back end processes events. Currently, there is no workaround.

BZ#2222683

Currently, there is no support for Multi-RHEL for the following deployment architectures:

- Edge (DCN)
- ShiftOnStack
- Director operator-based deployments
  Workaround: Use only a single version of RHEL across your RHOSP deployment when operating one of the listed architectures.

BZ#2223294

There is a known issue when performing an in-place upgrade from RHOSP 16.2 to 17.1 GA. The collection agent, **collectd-sensubility** fails to run on RHEL 8 Compute nodes.
Workaround: On affected nodes edit the file, **/var/lib/container-config-scripts/collectd_check_health.py**, and replace **"healthy: .State.Health.Status}"** with **"healthy: .State.Healthcheck.Status}"**/ on line 26.

BZ#2223916

In RHOSP 17.1 GA environments that use the ML2/OVN mechanism driver, there is a known issue with floating IP port forwarding not working correctly. This problem is caused because VLAN and flat networks distribute north-south network traffic when FIPs are used, and, instead, FIP port forwarding should be centralized on the Controller or the Networker nodes.

Workaround: To resolve this problem and force FIP port forwarding through the centralized gateway node, either set the RHOSP Orchestration service (heat) parameter **NeutronEnableDVR** to **false**, or use Geneve instead of VLAN or flat project networks.

BZ#2224236

In this release of 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#2224527

There is currently a known issue with the upgrade from RHOSP 16.2 to 17.1, when RADOS Gateway (RGW) is deployed as part of director-deployed Red Hat Ceph Storage. The procedure fails when HAProxy does not restart on the next stack update. Workaround: Apply the workaround from Red Hat KCS solution 7025985: HAProxy does not restart during RHOSP upgrade when RHCS is director-deployed and RGW is enabled

BZ#2225205

Outdated upgrade orchestration logic overrides the existing pacemaker authkey during the Fast Forward Upgrade (FFU) procedure, preventing Pacemaker from connecting to **pacemaker_remote** running on Compute nodes when Instance HA is enabled. As a result, the upgrade fails and **pacemaker_remote** running on Compute nodes is unreachable from the central cluster. Contact Red Hat support to receive instructions on how to perform FFU if Instance HA is configured.

BZ#2226366

There is currently a known issue when using a Red Hat Ceph Storage (RHCS) back end for volumes that can prevent instances from being rebooted, and may lead to data corruption. This occurs when all of the following conditions are met:

- RHCS is the back end for instance volumes.
- RHCS has multiple storage pools for volumes.
- A volume is being retyped where the new type requires the volume to be stored in a different pool than its current location.
- The retype call uses the **on-demand** migration_policy.
- The volume is attached to an instance.

Workaround: Do not retype **in-use** volumes that meet all of these listed conditions.

BZ#2227360

The image cache cleanup task of the NetApp NFS driver can cause unpredictable slowdowns in other Block Storage services. There is currently no workaround for this issue.

BZ#2229750

When you specify an availability zone (AZ) when creating a Block Storage volume backup, the AZ is

ignored. This may cause the backup to fail if the configuration of your AZs prevents the scheduler from satisfying the backup request. This issue does not affect the cross-availability-zone creation of volumes from existing backups.

BZ#2229761

There is currently a known issue with a race condition in the deployment steps for **ovn_controller** and **ovn_dbs**, which causes **ovn_dbs** to be upgraded before **ovn_controller**. If **ovn_controller** is not upgraded before **ovn_dbs**, an error before the restart to the new version causes packet loss. There is an estimated one-minute network outage if the race condition occurs during the Open Virtual Network (OVN) upgrade. A fix is expected in a later RHOSP release.

BZ#2229767

There is currently a known issue when you upgrade Red Hat Ceph Storage 4 to 5 during the upgrade from RHOSP 16.2 to 17.1. The **ceph-nfs** resource is misconfigured and Pacemaker does not manage the resource. The overcloud upgrade fails because the containers that are associated with **ceph-nfs-pacemaker** are down, impacting the Shared File Systems service (manila). A fix is expected in RHOSP 17.1.1. Workaround: Apply the workaround from Red Hat KCS solution 7028073: Pacemaker does not manage the **ceph-nfs** resource correctly during RHOSP and RHCS upgrade .

BZ#2229937

When **collectd sensubility** fails to create a sender, it does not close the link to the sender. Long-running open links that fail can cause issues in the bus, which cause **collectd sensubility** to stop working. Workaround: Restart the **collectd** container on affected overcloud nodes to recover **collectd sensubility**.

BZ#2231378

If you choose Red Hat Ceph Storage as the back end for your Block Storage (cinder) backup service repository, then you can only restore backed up volumes to a RBD-based Block Storage back end. There is currently no workaround for this.

BZ#2231893

The metadata service can become unavailable after the metadata agent fails in multiple attempts to start a malfunctioning HAProxy child container. The metadata agent logs an error message similar to: `ProcessExecutionError: Exit code: 125; Stdin: ; Stdout: Starting a new child container neutron-haproxy-ovnmeta-<uuid>".

Workaround: Run **podman kill <_container name_>** to stop the problematic haproxy child container.

BZ#2231960

When a Block Storage volume uses the Red Hat Ceph Storage back end, a volume cannot be removed when a snapshot is created from this volume and then a volume clone is created from this snapshot. In this case, you cannot remove the original volume while the volume clone exists.

BZ#2232171

If you download RHOSP 17.1.0 GA in the first few days of its availability, you might find that the version description in the file /etc/rhosp/release incorrectly includes the Beta designation, as shown in the following example.

```
(overcloud) [stack@undercloud-0 ~]$ cat /etc/rhosp-release
Red Hat OpenStack Platform release
17.1.0 Beta (Wallaby)
```

Workaround: If your GA deployment is affected, run the following command: **# dnf -y update rhosp-release**

BZ#2232199

If you download RHOSP 17.1.0 GA in the first few days of its availability, you might find that the version description in the file /etc/rhosp/release incorrectly includes the Beta designation, as shown in the following example.

```
(overcloud) [stack@undercloud-0 ~]$ cat /etc/rhosp-release
Red Hat OpenStack Platform release
17.1.0 Beta (Ussri)
```

Workaround: If your GA deployment is affected, run the following command: **# dnf -y update rhosp-release**

BZ#2233487

In RHOSP 17.1 GA environments that use RHOSP dynamic routing, there is a known issue where creating a load balancer using the RHOSP Load-balancing service with the OVN provider driver might fail. This failure can occur when there is latency between controller nodes. There is no workaround.

3.4.7. Deprecated functionality

The items in this section are either no longer supported, or will no longer be supported in a future release of Red Hat OpenStack Platform (RHOSP).

BZ#2128701

The ML2/OVS mechanism driver is deprecated since RHOSP 17.0.
Over several releases, Red Hat is replacing ML2/OVS with ML2/OVN. For instance, starting with RHOSP 15, ML2/OVN became the default mechanism driver.

Support is available for the deprecated ML2/OVS mechanism driver through the RHOSP 17 releases. During this time, the ML2/OVS driver remains in maintenance mode, receiving bug fixes and normal support, and most new feature development happens in the ML2/OVN mechanism driver.

In RHOSP 18.0, Red Hat plans to completely remove the ML2/OVS mechanism driver and stop supporting it.

If your existing RHOSP deployment uses the ML2/OVS mechanism driver, start now to evaluate a plan to migrate to the mechanism driver. Migration is supported in RHOSP 16.2 and 17.1.

Red Hat requires that you file a proactive support case before attempting a migration from ML2/OVS to ML2/OVN. Red Hat does not support migrations without the proactive support case. See How to open a proactive case for a planned activity on Red Hat OpenStack Platform? .

BZ#2136445

Monitoring of API health status via podman using sensubility is deprecated in RHOSP 17.1.
Only the sensubility layer is deprecated. API health checks remain in support. The sensubility layer exists for interfacing with Sensu, which is no longer a supported interface.

BZ#2139931

The metrics_qdr service (AMQ Interconnect) is deprecated in RHOSP 17.1. The metrics_qdr service continues to be supported in RHOSP 17.1 for data transport to Service Telemetry Framework (STF). The metrics_qdr service is used as a data transport for STF, and does not affect any other components for operation of Red Hat OpenStack.

BZ#2179428

Deploying the Block Storage (cinder) backup service in an active-passive configuration is

deprecated in RHOSP 17.1 and will be removed in a future release. For RHOSP 16.2 and RHOSP 17.0, the Block Storage (cinder) backup service is deployed in an active-passive configuration, and this configuration will continue to be supported in RHOSP 17.1 for these upgraded clusters.

BZ#2215264

Validations Framework (VF) is deprecated in RHOSP 17.1.

BZ#2238425

Collectd is deprecated in RHOSP 17.1.

3.4.8. Removed functionality

The items in this section are removed in this release of Red Hat OpenStack Platform (RHOSP):

BZ#2065541

In RHOSP 17.1, the collectd-gnocchi plugin is removed from director. You can use Service Telemetry Framework (STF) to collect monitoring data.

3.5. RED HAT OPENSTACK PLATFORM 17.1 BETA - JUNE 15, 2023

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

3.5.1. Bug fixes

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

BZ#1965308

Before this update, the Load-balancing service (octavia) could unplug a required subnet when you used different subnets from the same network as members' subnets. The members attached to this subnet were unreachable. With this update, the Load-balancing service does not unplug required subnets, and the load balancer can reach subnet members.

BZ#2066866

Even though the Panko monitoring service was deprecated, its endpoint still existed in the Identity service (keystone) after upgrading from RHOSP 16.2 to 17.1. With this update, the Panko service endpoint is cleaned up. However, Panko service users are not removed automatically. You must manually delete Panko service users with the command **openstack user delete panko**. There is no impact if you do not delete these users.

BZ#2080199

Before this update, services that were removed from the undercloud were not cleaned up during upgrades from RHOSP 16.2 to 17.0. The removed services remained in the OpenStack endpoint list even though they were not reachable or running. With this update, RHOSP upgrades include Ansible tasks to clean up the endpoints that are no longer required.

BZ#2097844

Before this update, the **overcloud config download** command failed with a traceback error because the command attempted to reach the Orchestration service (heat) to perform the download. The Orchestration service is no longer persistently running on the undercloud. With this update, the **overcloud config download** command is removed. Instead, you can use your **overcloud deploy** command with the **--stack-only** option.

BZ#2116600

Sometimes, during a live migration, a libvirt internal error **migration was active, but no RAM info was set** was raised even though the live migration was successful. The live migration failed when it

should have succeeded. With this update, when this libvirt internal error is raised, the live migration is signaled as complete in the libvirt driver. The live migration correctly succeeds in this condition.

BZ#2125610

Before this update, an SELinux issue triggered errors with RHOSP Load-balancing service (octavia) ICMP health monitors that used the Amphora provider driver. In RHOSP 17.1, this issue has been fixed and ICMP health monitors function correctly.

BZ#2125612

Before this update, users might have experienced the following warning message in the Load-balancing service (octavia) Amphora VM log file when the load balancer was loaded with multiple concurrent sessions: **nf_conntrack: table full, dropping packet**. This error occurred if the Amphora VM dropped Transport Control Protocol (TCP) flows and caused latency on user traffic. With this update, connection tracking (conntrack) is disabled for TCP flows in the Load-balancing service Amphora VM, and new TCP flows are not dropped. Conntrack is only required for User Datagram Protocol (UDP) flows.

BZ#2129207

Before this update, a network disruption or temporary unavailability of the Identity service (keystone) resulted in the nova-conductor service failing to start. With this update, the nova-conductor service logs a warning and continues startup in the presence of disruptions that are likely to be temporary. As a result, the nova-conductor service does not fail to start if transient issues like network disruptions or temporary unavailability of necessary services are encountered during startup.

BZ#2133027

The Alarming service (aodh) uses the deprecated gnocchi API to aggregate metrics, which results in incorrect metric measures of CPU usage in gnocchi. With this update, dynamic aggregation in gnocchi supports the ability to make re-aggregations of existing metrics and the ability to manipulate and transform metrics as required. CPU time in gnocchi is correctly calculated.

BZ#2133297

Before this update, the **openstack undercloud install** command launched the **openstack tripleo deploy** command, which created the **/home/stack/.tripleo/history** file with **root:root** as the owner. Subsequent deploy commands failed because of permission errors. With this update, the command creates the file with the **stack** user as the owner, and deploy commands succeed without permission errors.

BZ#2140988

Before this update, a live migration might fail because the database did not update with the destination host details.
With this update, the instance host value in the database is set to the destination host during live migration.

BZ#2149216

Before this update Open Virtual Network (OVN) load balancer health checks were not performed if you used Floating IPs (FIP) associated to the Load Balancer Virtual IP (VIP), and traffic was redirected to members in the Error state if the FIP was used.
With this update, if you use Floating IPs (FIP) is associated to the Load Balancer Virtual IP (VIP), there is a new load balancer health check created for the FIP, and traffic is not redirected to members in the Error state.

BZ#2149468

Before this update, the Compute service (nova) processed a temporary error message from the Block Storage service (cinder) volume detach API, such as '504 Gateway Timeout', as an error. The Compute service failed the volume detach operation even though it succeeded but timed out on the

Block Storage service side, leaving a stale block device mapping record in the Compute service database. With this update, the Compute service retries the volume detach call to the Block Storage service API if it receives an HTTP error that is likely to be temporary. Upon retry, if the volume attachment is no longer found, the Compute service processes the volume as already detached.

BZ#2151043

Before this update, the **openstack-cinder-volume-0** container, which is created by the Pacemaker bundle resource for the Block Storage service (cinder), mounted /**run** from the host. This mount path created the **.containerenv** file in the directory. When the **.containerenv** file exists, **subscription-manager** fails because it evaluates that the command is executed inside a container. With this update, the mount path is updated so that Podman disables the creation of the **.containerenv** file, and **subscription-manager** executes successfully in a host that is running the **openstack-cinder-volume-0** container.

BZ#2152888

Before this update, the Service Telemetry Framework (STF) API health monitoring script was failing because it depended on Podman log content, which was no longer available. With this update, the health monitoring script depends on the Podman socket instead of the Podman log, and API health monitoring operates normally.

BZ#2154343

Before this update, the disabling and enabling of network log objects in a security group was inconsistent. The logging of a connection was disabled as soon as one of the log objects in the security group associated with that connection was disabled. With this update, a connection is logged if any of the related enabled log objects in the security group allow it, even if one of those log objects becomes disabled.

BZ#2162756

Before this update, VLAN network traffic was centralized over the Controller nodes. With this update, if all the tenant provider networks that are connected to a router are of the VLAN/Flat type, that traffic is now distributed. The node that contains the VM sends the traffic directly.

BZ#2163815

Before this update, Open Virtual Network (OVN) load balancers on switches with **localnet** ports (Networking service (neutron) provider networks) did not work if traffic came from **localnet**. With this update, load balancers are not added to the logical switch associated with the provider network. This update forces Network Address Translation (NAT) to occur at the virtual router level instead of the logical switch level.

BZ#2164421

Before this update, the Compute service (nova) did not confidence-check the content of Virtual Machine Disk (VMDK) image files. By using a specially crafted VMDK image, it was possible to expose sensitive files on the host file system to guests booted with that VMDK image. With this update, the Compute service confidence checks VMDK files and forbids VMDK features that the leak behavior depends on. It is no longer possible to leak sensitive host file system contents using specially crafted VMDK files.

NOTE

Red Hat does not support the VMDK image file format in RHOSP.

BZ#2164677

Before this update, the iptables rule for the heat-cfn service contained the incorrect TCP port number. Users could not access the heat-cfn service endpoint if SSL was enabled for public endpoints. With this update, the TCP port number is correct in the iptables rule. Users can access the heat-cfn service endpoint, even if SSL is enabled for public endpoints.

BZ#2167161

Before this update, the default value of **rgw_max_attr_size** was 256, which created issues for OpenShift on OpenStack when uploading large images. With this update, the default value of **rgw_max_attr_size** is 1024.
You can change the value by adding the following configuration to an environment file that you include in your overcloud deployment:

```
parameters_default:
  CephConfigOverrides:
    rgw_max_attr_size: <new value>
```

BZ#2169303

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. The IPMI agent container is executed with the **--privilege** flag to execute **ipmitool** commands on the host. The Telemetry service (ceilometer) can now capture power metrics.

BZ#2169349

Before this update, instances were losing communication with the ovn-metadata-port because the load balancer health monitor was replying to the ARP requests for the OVN metadata agent's IP, causing the request going to the metadata agent to be sent to another MAC address. With this update, the ovn-controller conducts back-end checks by using a dedicated port instead of the ovn-metadata-port. When establishing a health monitor for a load balancer pool, ensure that there is an available IP in the VIP load balancer's subnet. This port is distinct for each subnet, and various health monitors in the same subnet can reuse the port. Health monitor checks no longer impact ovn-metadata-port communications for instances.

BZ#2172063

Before this update, the **openstack overcloud ceph deploy** command may have failed during the **apply spec** operation if the chrony NTP service was down. With this update, the chrony NTP service is enabled before the **apply spec** operation.

BZ#2172582

Before this update, the **create pool** operation failed because the podman command used **/etc/ceph** as the volume argument. This argument does not work for Red Hat Ceph Storage version 6 containers. With this update, the podman command uses **/var/lib/ceph/$FSID/config/** as the first volume argument and **create pool** operations are successful.

BZ#2173101

Before this update, when users deployed Red Hat Ceph Storage in a tripleo-ipa context, a **stray hosts** warning showed in the cluster for the Ceph Object Gateway (RADOS Gateway [RGW]). With this update, during a Ceph Storage deployment, you can pass the option **--tld** in a tripleo-ipa context to use the correct hosts when you create the cluster.

BZ#2173575

Before this update, when a VM that was associated to a provider network with disabled port security attempted to reach IPs on the provider network that were not recognized by OpenStack, there was a flooding issue because the forwarding database (FDB) table was not learning MAC addresses. This patch uses a new option in OVN to enable the learning of IPs in the FDB table. There is currently no ageing mechanism for the FDB table. You can clean up the table periodically to prevent the occurrence of scaling issues caused by the size of the table.

BZ#2178618

Before this update, a security group logging enhancement introduced an issue where log objects could not be deleted at the same time as security groups. This action caused an internal server error. With this update, the **db_set** function that modifies the northbound database entries does not fail if

the row that is requested does not exist any more.

BZ#2180933

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/**. Because of this issue, the rsyslog service could not locate Pacemaker log files. With this update, the Pacemaker log path is changed from **/var/log/pacemaker/** to **/var/log/host/pacemaker/**.

BZ#2188252

Before this update, the 'openstack tripleo container image prepare' command failed because there were incorrect Ceph container tags in the **container_image_prepare_defaults.yaml** file. With this update, the correct Ceph container tags are in the YAML file, and the 'openstack tripleo container image prepare' command is successful.

BZ#2203238

Before this update, for the nova-compute log to record os-brick privileged commands for debugging purposes, you had to apply the workaround outlined in https://access.redhat.com/articles/5906971. This update makes the workaround redundant and provides a better solution that separates logging by the nova-compute service so that the privileged commands of os-brick are logged at the debug level but the privileged commands of nova are not.

3.5.2. Enhancements

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

BZ#1369007

Cloud users can launch instances that are protected with UEFI Secure Boot when the overcloud contains UEFI Secure Boot Compute nodes. For information on creating an image for UEFI Secure Boot, see Creating an image for UEFI Secure Boot . For information on creating a flavor for UEFI Secure Boot, see "UEFI Secure Boot" in Flavor metadata.

BZ#1581414

Before this release, **NovaHWMachineType** could not be changed for the lifetime of a RHOSP deployment because the machine type of instances without a **hw_machine_type** image property would use the newly configured machine types after a hard reboot or migration. Changing the underlying machine type for an instance could break the internal ABI of the instance.
With this release, when launching an instance the Compute service records the instance machine type within the system metadata of the instance. Therefore, it is now possible to change the **NovaHWMachineType** during the lifetime of a RHOSP deployment without affecting the machine type of existing instances.

BZ#1619266

This update introduces the security group logging feature. To monitor traffic flows and attempts into and out of a virtual machine instance, you can configure the Networking Service packet logging for security groups.
You can associate any virtual machine instance port with one or more security groups and define one or more rules for each security group. For instance, you can create a rule to drop inbound ssh traffic to any virtual machine in the finance security group. You can create another rule to allow virtual machines 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.

You can use security group logging for both stateful and stateless security groups.

Logged events are stored on the compute nodes that host the virtual machine instances, in the file **/var/log/containers/stdouts/ovn_controller.log**.

BZ#1672972

This enhancement helps cloud users determine if the reason they are unable to access an "ACTIVE" instance is because the Compute node that hosts the instance is unreachable. RHOSP administrators can now configure the following parameters to enable a custom policy that provides a status in the **host_status** field to cloud users when they run the **openstack show server details** command, if the host Compute node is unreachable:

- **NovaApiHostStatusPolicy**: Specifies the role the custom policy applies to.
- **NovaShowHostStatus**: Specifies the level of host status to show to the cloud user, for example, "UNKNOWN".

BZ#1693377

With this update, an instance can have a mix of shared (floating) CPUs and dedicated (pinned) CPUs instead of only one CPU type. RHOSP administrators can use the **hw:cpu_policy=mixed** and **hw_cpu_dedicated_mask** flavor extra specs to create a flavor for instances that require a mix of shared CPUs and dedicated CPUs.

BZ#1701281

In RHOSP 17.1, support is available for cold migrating and resizing instances that have vGPUs.

BZ#1761861

With this update, you can configure each physical GPU on a Compute node to support a different virtual GPU type.

BZ#1761903

On RHOSP deployments that use a routed provider network, you can now configure the Compute scheduler to filter Compute nodes that have affinity with routed network segments, and verify the network in placement before scheduling an instance on a Compute node. You can enable this feature by using the **NovaSchedulerQueryPlacementForRoutedNetworkAggregates** parameter.

BZ#1772124

With this update, you can use the new **NovaMaxDiskDevicesToAttach** heat parameter to specify the maximum number of disk devices that can be attached to a single instance. The default is unlimited (-1). For more information, see Configuring the maximum number of storage devices to attach to one instance.

BZ#1782128

In RHOSP 17.1, a RHOSP administrator can provide cloud users the ability to create instances that have emulated virtual Trusted Platform Module (vTPM) devices. RHOSP only supports TPM version **2.0**.

BZ#1793700

In RHOSP 17.1, a RHOSP administrator can declare which custom physical features and consumable resources are available on the RHOSP overcloud nodes by modeling custom traits and inventories in a YAML file, **provider.yaml**.

BZ#1827598

This RHOSP release introduces support of the OpenStack stateless security groups API.

BZ#1873409

On RHOSP deployments that are configured for OVS hardware offload and to use ML2/OVN, and that have Compute nodes with VDPA devices and drivers and Mellanox NICs, you can enable your cloud users to create instances that use VirtIO data path acceleration (VDPA) ports. For more

information, see Configuring VDPA Compute nodes to enable instances that use VDPA ports and Creating an instance with a VDPA interface .

BZ#1873707

With this update, you can use the validation framework in the workflow of backup and restore procedures to verify the status of the restored system. The following validations are included:

- **undercloud-service-status**
- **neutron-sanity-check**
- **healthcheck-service-status**
- **nova-status**
- **ceph-health**
- **check-cpu**
- **service-status**
- **image-serve**
- **pacemaker-status**
- **validate-selinux**
- **container-status**

BZ#1883554

With this update, a RHOSP administrator can now create a flavor that has a **socket** PCI NUMA affinity policy, which can be used to create an instance that requests a PCI device only when at least one of the instance NUMA nodes has affinity with a NUMA node in the same host socket as the PCI device.

BZ#1962500

With this update, you can configure the collectd logging source in TripleO Heat Templates. The default value matches the default logging path.

BZ#2033811

The Shared File System service (manila) now supports using Pure Storage Flashblade system as a backend. Refer to the Red Hat ecosystem catalog to find the vendor's certification and installation documentation.

BZ#2066349

With this enhancement, the LVM volumes installed by the **overcloud-hardened-uefi-full.qcow2** whole disk overcloud image are now backed by a thin pool. The volumes are still grown to consume the available physical storage, but are not over-provisioned by default.
The benefits of thin-provisioned logical volumes:

- If a volume fills to capacity, the options for manual intervention now include growing the volume to over-provision the physical storage capacity.
- The RHOSP upgrades process can now create ephemeral backup volumes in thin-provisioned environments.

BZ#2069624

The RHOSP snapshot and revert feature is based on the Logical Volume Manager (LVM) snapshot functionality and is intended to revert an unsuccessful upgrade or update. Snapshots preserve the original disk state of your RHOSP cluster before performing an upgrade or an update. You can then remove or revert the snapshots depending on the results. If an upgrade completed successfully and you do not need the snapshots anymore, remove them from your nodes. If an upgrade fails, you can revert the snapshots, assess any errors, and start the upgrade procedure again. A revert leaves the disks of all the nodes exactly as they were when the snapshot was taken.

BZ#2104522

With this update, live migration now uses multichassis Open Virtual Network (OVN) ports to optimize the migration procedure and significantly reduce network downtime for VMs during migration in particular scenarios.

BZ#2111528

With this update, the default Ceph container image is based on Red Hat Ceph Storage 6 instead of Red Hat Ceph Storage 5.

BZ#2124309

With this enhancement, operators can enable the run_arping feature for Pacemaker-managed virtual IPs (VIPs), so that the cluster preemptively checks for duplicate IPs. To do this, you must add the following configuration to the environment file: ExtraConfig: pacemaker::resource::ip::run_arping: true If a duplicate is found, the following error is logged in the /var/log/pacemaker/pacemaker.log log file: Sep 07 05:54:54 IPaddr2(ip-172.17.3.115)[209771]: ERROR: IPv4 address collision 172.17.3.115 [DAD] Sep 07 05:54:54 IPaddr2(ip-172.17.3.115)[209771]: ERROR: Failed to add 172.17.3.115

BZ#2133055, BZ#2138238

With this update, you deploy two separate instances of the Image service (glance) API. The instance that is accessible to OpenStack tenants is configured to hide image location details, such as the direct URL of an image or whether the image is available in multiple locations. The second instance is accessible to OpenStack administrators and OpenStack services, such as the Block Storage service (cinder) and the Compute service (nova). This instance is configured to provide image location details. This enhancement addresses the recommendations of OSSN-0090 and CVE-2022-4134. With this update, a malicious user cannot leverage the location details of an image to upload an altered image.

BZ#2152877

This enhancement adds OVN security group logging to the Networking service (neutron) for the reply packets of a network connection. The ovn-controller log files now log the full network connection.

BZ#2165501

Starting with Red Hat OpenStack Platform (RHOSP) 17.1, in ML2/OVN deployments, you can enable hardware offloading on minimum bandwidth or bandwidth limit QoS egress policies. You cannot enable hardware offloading on ingress policies. For more information, see Configuring the Networking service for QoS policies.

3.5.3. Technology previews

The items listed in this section are provided as Technology Previews for 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#1813561

With this update, the Load-balancing service (octavia) supports HTTP/2 load balancing by using the Application Layer Protocol Negotiation (ALPN) for listeners and pools that are enabled with Transport Layer Security (TLS). The HTTP/2 protocol improves performance by loading pages faster.

BZ#1848407

In RHOSP 17.1, a technology preview is available for the Stream Control Transmission Protocol (SCTP) in the Load-balancing service (octavia). Users can create SCTP listeners and attach SCTP pools in a load balancer.

BZ#2057921

In RHOSP 17.1, a technology preview is available for creating load balancers over an IPv6 management network. Using a private IPv6 management network for the Load-balancing service (octavia) may simplify edge deployments.

BZ#2088291

In RHOSP 17.1, a technology preview is available for ML2/OVN QoS bandwidth limiting for router gateway IP ingress and egress.

3.5.4. Release notes

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

BZ#2178015

In RHOSP 17.1, Red Hat recommends that all physical functions (PFs) on the same NIC hardware use drivers that are in the same space. PFs on the same NIC should all use drivers that run in either the user space or in the kernel space.
For example, if PF1 on NIC1 is used by the DPDK PMD driver, then PF2 on NIC1 should not use the kernel driver. In this example, the PFs on NIC1 should both use the DPDK PMD driver or both use the kernel driver.

3.5.5. Known issues

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

BZ#2108212

If you use IPv6 to connect to VM instances during migration to the OVN mechanism driver, connection to the instances might be disrupted for up to several minutes when the ML2/OVN services start.
The router advertisement daemon **radvd** for IPv6 is stopped during migration to the OVN mechanism driver. While **radvd** is stopped, router advertisements are no longer broadcast. This broadcast interruption results in VM instance connection loss over IPv6. IPv6 communication is automatically restored once the new ML2/OVN services start.

Workaround: To avoid the potential disruption, use IPv4 instead.

BZ#2109985

Currently, in ML2/OVS deployments, Open vSwitch (OVS) does not support offloading OpenFlow rules that have the **skb_priority**, **skb_mark**, or output queue fields set. These fields are required for Quality of Service (QoS) support for virtio ports.
If you set a minimum bandwidth rule for a virtio port, the Networking service (neutron) OVS agent marks the traffic of this port with a Packet Mark field. 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 that no traffic can be offloaded.

Workaround: If your environment includes OVS hardware offload ports, disable packet marking in the

nodes that require hardware offloading. When you disable packet marking, it is not possible to set rate limiting rules for virtio ports. However, differentiated services code point (DSCP) marking rules are still available.

In the configuration file, set the **disable_packet_marking** flag to **true**. When 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#2126810

In RHOSP 17.0, the DNS service (designate) and the Load-balancing service (octavia) are misconfigured for high availability. The RHOSP Orchestration service (heat) templates for these services use the non-Pacemaker version of the Redis template.
Workaround: include **environments/ha-redis.yaml** in the **overcloud deploy** command after the **enable-designate.yaml** and **octavia.yaml** environment files.

BZ#2144492

If you migrate a RHOSP 17.1 ML2/OVS deployment with centralized routing (no DVR) to ML2/OVN, the floating IP (FIP) downtime that occurs during ML2/OVN migration can exceed 60 seconds.

BZ#2160481

In RHOSP 17.1 environments that use BGP dynamic routing, there is currently a known issue where floating IP (FIP) port forwarding fails.
When FIP port forwarding is configured, packets sent to a specific destination port with a destination IP that equals the FIP are redirected to an internal IP from a RHOSP Networking service (neutron) port. This occurs regardless of the protocol that is used: TCP, UDP, and so on.

When BGP dynamic routing is configured, the routes to the FIPs used to perform FIP port forwarding are not exposed, and these packets cannot reach their final destinations.

Currently, there is no workaround.

BZ#2163477

In RHOSP 17.1 environments that use BGP dynamic routing, there is currently a known issue affecting VM instances connected to provider networks. The RHOSP Compute service cannot route packets sent from one of these VM instances to a multicast IP address destination. Therefore, VM instances subscribed to a multicast group fail to receive the packets sent to them. The cause is that BGP multicast routing is not properly configured on the overcloud nodes. Currently, there is no workaround.

BZ#2182371

There is currently a known issue with guest instances that use Mellanox ConnectX-5, ConnectX-6, and Bluefield-2 NICs with offload (switchdev) ports. It takes a long time to initialize the system when you reboot the operating system from the guest directly, for example, by using the command **sudo systemctl reboot --reboot-arg=now**. If the VM is configured with two Virtual Functions (VFs) from the same Physical Function (PF), the initialization of one of the VFs might fail and cause a longer initialization time. Workaround: Reboot the guest instance in a timely manner by using the OpenStack API instead of rebooting the guest instance directly.

BZ#2183793

Red Hat has not validated the RHOSP 17.1 beta release on NFV deployments with AMD processors. Testing is underway now with plans to validate the application in a future release.

Do not use RHOSP 17.1 NFV deployments with AMD hardware for production until Red Hat validates the application. Any use of this pre-tested application is at risk for unintended results.

BZ#2184070

This update adds a check to ensure that there are enough IP addresses available for each subnet pool during an OVN migration. If you do not have enough IP addresses, the migration script will stop and display a warning.

BZ#2185897

In ML2/OVN deployments, do not use live migration on virtual machine instances that use trunk ports. On instances that use trunk ports, live migration can fail due to the flapping of the instance's subport between the compute nodes. For instances that have trunk ports, use cold migration instead.

BZ#2192913

In RHOSP 17.1 environments with ML2/OVN, DVR enabled and using VLAN tenant networks, east/west traffic between VMs connected to different tenant networks is flooded to the fabric. As a result, packets between those VMs reach not only the compute nodes where those VMs run, but also any other overcloud node.

This could cause an impact in the network side and it could be a security risk because the fabric sends traffic everywhere.

This bug will be fixed in a later FDP release, so no RHOSP update is needed to obtain it.

BZ#2196291

There is currently a known issue wherein custom SRBAC rules do not permit list policy rules to non-admin users. As a consequence, non-admin users can not list or manage these rules. Current workarounds include either disabling SRBAC, or modifying the SRBAC custom rule to permit this action.

BZ#2203857

Currently, a known issue in the Ceph RADOS Gateway component in Red Hat Ceph Storage (RHCS) 6.0 causes authorization with Identity service (keystone) tokens to fail. See https://bugzilla.redhat.com/2188266.
As a result, when you configure your deployment with Red Hat Ceph Storage using RADOS Gateway as the object-store server, Object Storage service (swift) clients fail and return code **403/Unauthorized**. The issue did not manifest in tests that deployed pre-release versions of RHCS 6.1, which was released for general availability on June 15, 2023.

Also, OpenShift integration on OpenStack has not been validated for beta because the default configuration uses RADOS Gateway. The following workaround is expected to mitigate the issue and enable you to do preliminary tests with OpenShift integration on OpenStack.

Workaround: Deploy the Object Storage service (swift) as the object-store server instead of RADOS Gateway, even when enabling Ceph Storage for persistent Block Storage service (cinder) or Image service (glance) storage and ephemeral Compute service (nova) storage. To do this, replace the **cephadm.yaml** environment file with the **cephadm-rbd-only.yaml** in the deployment command line.

When you configure the OpenStack environment with the Object Storage service (swift) instead of RADOS Gateway as the object-store server, Object Storage service (swift) clients work as expected.

BZ#2207991

Currently, secure role-based access control (SRBAC) and the **NovaShowHostStatus** parameter use the same policy key titles. If you configure both SRBAC and **NovaShowHostStatus**, the deployment

fails with a conflict. In RHOSP 17.1-Beta, you cannot use both features in the same deployment. A fix is expected in the RHOSP 17.1 GA release.

BZ#2210030

There is currently a known issue where custom SRBAC rules do not permit list shared security groups to non-administrative users that are not rule owners. This causes shared security groups and rules to not be managed properly by non-administrative users that are not rule owners. Workaround: Disable custom SRBAC rules or modify the custom rules to permit any user to manage the rules.

BZ#2210062

In RHOSP 17.1 environments that use BGP dynamic routing with OVN, there is a known issue where the default value of the Autonomous System Number (ASN) used by the OVN BGP agent differs from the ASN used by FRRouting (FRR).
Workaround: ensure that the values for the tripleo parameters used in the undercloud and overcloud configuration, **FrrBgpAsn** and **FrrOvnBgpAgentAsn**, are identical.

BZ#2210319

There is currently a known issue where the Retbleed vulnerability mitigation in RHEL 9.2 can cause a performance drop for Open vSwitch with Data Plane Development Kit (OVS-DPDK) on Intel Skylake CPUs.
This performance regression happens only if C-states are disabled in the BIOS, hyper-threading is enabled, and OVS-DPDK is using only one hyper-thread of a given core.

Workaround: Assign both hyper-threads of a core to OVS-DPDK or to SRIOV guests that have DPDK running as recommended in the NFV configuration guide.

BZ#2211691

There is currently a known issue where changes to the Block Storage service (cinder), related to CVE-2023-2088, impact the ability of the Bare Metal Provisioning service (ironic) to detach a volume that is attached to a physical bare metal node. The detachment is required for the teardown of physical machines with an instance deployed on them. You can deploy bare-metal instances by using the Compute service (nova) or by using the boot from volume functionality. However, you cannot automatically tear down instances by using boot from Block Storage service volumes. There is no workaround for this issue. A fix is expected in the RHOSP 17.1 GA release.

BZ#2211849

In RHOSP 17.1 environments that use BGP dynamic routing, there is currently a known issue where the OVN BGP agents that are running on overcloud nodes fail because of a bug in a shipped library (pyroute2). When this issue occurs, no new routes are advertised from the affected node, and there might be a loss of connectivity with new or migrated VMs, new load balancers, and so on.
Workaround: Install an updated version of pyroute2 in the **ovn_bgp_agent** container, by adding the following lines to **containers-prepare-parameter.yaml**:

```
ContainerImagePrepare:
- push_destination: true
  ...
  includes:
  - nova-compute
  modify_role: tripleo-modify-image
  modify_append_tag: "-hotfix"
  modify_vars:
    tasks_from: rpm_install.yml
    rpms_path: /home/stack/nova-hotfix-pkgs
  ...
```

For more information, see Installing additional RPM files to container images .

BZ#2213126

The logging queue that buffers excess security group log entries sometimes stops accepting entries before the specified limit is reached. As a workaround, you can set the queue length higher than the number of entries you want it to hold.
You can set the maximum number of log entries per second with the parameter **NeutronOVNLoggingRateLimit**. If the log entry creation exceeds that rate, the excess is buffered in a queue up to the number of log entries that you specify in **NeutronOVNLoggingBurstLimit**.

The issue is especially evident in the first second of a burst. In longer bursts, such as 60 seconds, the rate limit is more influential and compensates for burst limit inaccuracy. Thus, the issue has the greatest proportional effect in short bursts.

Workaround: Set **NeutronOVNLoggingBurstLimit** at a higher value than the target value. Observe and adjust as needed.

BZ#2214328

Currently, DNS-as-a-Service (designate) is misconfigured when secure role-based access control (SRBAC) is enabled. If you configure both SRBAC and DNS-as-a-Service, the RHOSP deployment fails. Workaround: For a successful deployment, apply the following patches on the undercloud server:

- https://review.opendev.org/c/openstack/tripleo-heat-templates/+/885291
- https://review.opendev.org/c/openstack/tripleo-heat-templates/+/885599

BZ#2215053

In RHOSP 17.1 environments that use Border Gateway Protocol (BGP) dynamic routing, there is currently a known issue where the FRRouting (FRR) container fails to deploy. This failure occurs because the RHOSP director deploys the FRR container before the container image prepare task finishes. Workaround: In your heat templates, ensure that the **ContainerImagePrepare** precedes the **overcloud deploy** command.

BZ#2215936

If you migrate from ML2/OVS with SR-IOV to ML2/OVN, and then attempt to create a VM instance with virtual functions (VF), the instance creation fails. The problem does not affect instances with physical functions (PF).

3.5.6. Deprecated functionality

The items in this section are either no longer supported, or will no longer be supported in a future release of Red Hat OpenStack Platform (RHOSP):

BZ#2128701

The ML2/OVS mechanism driver is deprecated since RHOSP 17.0.
Over several releases, Red Hat is replacing ML2/OVS with ML2/OVN. For instance, starting with RHOSP 15, ML2/OVN became the default mechanism driver.

Support is available for the deprecated ML2/OVS mechanism driver through the RHOSP 17 releases. During this time, the ML2/OVS driver remains in maintenance mode, receiving bug fixes and normal support, and most new feature development happens in the ML2/OVN mechanism driver.

In RHOSP 18.0, Red Hat plans to completely remove the ML2/OVS mechanism driver and stop supporting it.

If your existing RHOSP deployment uses the ML2/OVS mechanism driver, start now to evaluate a plan to migrate to the mechanism driver. Migration is supported in RHOSP 16.2 and 17.1.

Red Hat requires that you file a proactive support case before attempting a migration from ML2/OVS to ML2/OVN. Red Hat does not support migrations without the proactive support case. See How to submit a Proactive Case.

BZ#2136445

Monitoring of API health status via podman using sensubility is deprecated in RHOSP 17.1.
Only the sensubility layer is deprecated. API health checks remain in support. The sensubility layer exists for interfacing with Sensu, which is no longer a supported interface.

BZ#2139931

The metrics_qdr service (AMQ Interconnect) is deprecated in RHOSP 17.1. The metrics_qdr service continues to be supported in RHOSP 17.1 for data transport to Service Telemetry Framework (STF). The metrics_qdr service is used as a data transport for STF, and does not affect any other components for operation of Red Hat OpenStack.

BZ#2179428

Deploying the Block Storage (cinder) backup service in an active-passive configuration is deprecated in RHOSP 17.1 and will be removed in a future release. For RHOSP 16.2 and RHOSP 17.0, the Block Storage (cinder) backup service is deployed in an active-passive configuration, and this configuration will continue to be supported in RHOSP 17.1 for these upgraded clusters.

BZ#2215264

Validations Framework (VF) is deprecated in RHOSP 17.1.

3.5.7. Removed functionality

The items in this section are removed in this release of Red Hat OpenStack Platform (RHOSP):

BZ#2065541

In RHOSP 17.1, the collectd-gnocchi plugin is removed from director. You can use Service Telemetry Framework (STF) to collect monitoring data.

BZ#2126890

The Derived Parameters feature is removed. The Derived Parameters feature is configured using the **--plan-environment-file** option of the **openstack overcloud deploy command**.
Workaround / Migration Instructions

NFV and HCI overclouds require system tuning. There are many different options for system tuning. The Derived Parameters functionality tuned systems with director using to inspect hardware inspection data and set tuning parameters using the **--plan-environment-file** option of the **openstack overcloud deploy** command. The Derived Parameters functionality is removed in 17.1.

The following parameters were tuned by this functionality:

- IsolCpusList
- KernelArgs
- NeutronPhysnetNUMANodesMapping

- NeutronTunnelNUMANodes
- NovaCPUAllocationRatio
- NovaComputeCpuDedicatedSet
- NovaComputeCpuSharedSet
- NovaReservedHostMemory
- NovaReservedHostMemory
- OvsDpdkCoreList
- OvsDpdkSocketMemory
- OvsPmdCoreList
  To set and tune these parameters, observe their values using the available command line tools and set them using a standard heat template.

CHAPTER 4. DOCUMENTATION CHANGES

This section details the major documentation updates delivered with Red Hat OpenStack Platform (RHOSP) 17.1, and the changes made to the documentation set that include adding new features, enhancements, and corrections. The section also details the addition of new titles and the removal of retired or replaced titles.

Table 4.1. Documentation changes legend

Column	Meaning
Date	The date that the documentation change was published.
17.1 versions impacted	The RHOSP 17.1 versions that the documentation change impacts. Unless stated otherwise, a change that impacts a particular version also impacts all later versions.
Components	The RHOSP components that the documentation change impacts.
Affected content	The RHOSP documents that contain the change or update.
Description of change	A brief summary of the change to the document.

Table 4.2. Document changes

Date	17.1 versions impacted	Components	Affected content	Description of change
06 September 2024	17.1.2	Networking	Migrating to the OVN mechanism driver, ML2/OVS to ML2/OVN in-place migration: validated and prohibited scenarios.	Removed limitation. You can now migrate to the OVN mechanism driver if the original ML2/OVS environment uses iptables_hybrid firewall and trunk ports.
04 September 2024	17.1.3	Security	Troubleshooting Active Directory integration, Troubleshooting Red Hat Identity (IdM) integration.	Corrected LDAP search commands.
22 August 2024	17.1.3	Security	Migrating instances	Removed section 'Disable live migration', please obtain a support exception if you want to disable live migration.

Date	17.1 versions impacted	Components	Affected content	Description of change
20 August 2024	17.1.3	Bare Metal Provisioning	Enabling ISO boot for bare-metal instances Booting an ISO image directly for use as a RAM disk	Added procedures that detail how to boot an ISO image directly for use as a RAM disk.
03 July 2024	17.1	Performance and Scaling	Tuning the undercloud	Added recommendation to increase message sizes.
05 June 2024	17.1	Edge	Migrating to a multistack deployment	Removed statement that migrating from single stack to multistack is unsupported.
23 May 2024	17.1	Networking	Layer 3 high availability with OVN	Added a note prohibiting the use of the **--ha** option when creating an OVN router.
21 May 2024	17.1	Networking	Changing Load-balancing service default settings	Documented support for IPv6 for the Load-balancing service (octavia) amphora control subnet.
15 May 2024	17.1	Networking	Creating custom virtual routers with router flavors	Documented the new Networking service (neutron) plug-in, **ovn-router-flavors-ha**.
15 May 2024	17.1	NFV	Saving power in OVS-DPDK deployments	Documented new power saving profile.
29 April 2024	17.1	Security	Static Media	Removed unsupported content
29 March 2024	17.1	Networking	Monitoring OVN database status.	A new section describes how to monitor OVN databases.

Date	17.1 versions impacted	Components	Affected content	Description of change
22 March 2024	17.1	NFV	• Chapter 8. Configuring OVS TC-flower hardware offload • Chapter 7. Configuring an SR-IOV deployment • Chapter 10. Configuring an OVS-DPDK deployment	There is a new chapter written about OVS TC-flower hardware. Chapters on SR-IOV and OVS-DPDK have been rewritten.
21 March 2024	17.1	Compute, Networking	Tagging virtual devices	Moved this section from the the *Configuring Red Hat OpenStack Platform Networking* guide to the *Creating and managing instances* guide. The commands have been updated and the new content includes how to tag both block devices and virtual NICs while attaching them to an existing instance.
14 March 2024	17.1	Storage	Configuring NFS storage	Added this section to the Block Storage service back ends topic in *Configuring persistent storage*.
13 March 2024	17.1	Storage	Manage and unmanage volumes and their snapshots	Added a new topic that describes the reasons and the associated commands, for managing and unmanaging Block Storage volumes and their snapshots.
4 March 2024	17.1	Networking	Replacing a bootstrap Controller node	Now you can use the original hostname and IP address for the replacement Controller node when you replace a Controller node.

Date	17.1 versions impacted	Components	Affected content	Description of change
29 February 2024	17.1	Deployment	Installing and managing Red Hat OpenStack Platform with director Customizing your Red Hat OpenStack Platform deployment	Updated the content in *Installing and managing Red Hat OpenStack Platform with director* guide to focus only on the core tasks required to deploy a basic RHOSP environment. Content related to optional features and custom configuration are moved into a new guide: *Customizing your Red Hat OpenStack Platform deployment*.
27 February 2024	17.1	NFV	Example Ceph configuration file	The topic, "Example Ceph configuration file," has been updated.
19 February 2024	17.1	Networking	• Configuring floating IP port forwarding • Creating port forwarding for a floating IP	Two new topics about floating IP port forwarding have been added to *Configuring Red Hat OpenStack Platform networking*.
16 February 2024	17.1	Security	Using Fernet keys for encryption in the overcloud	Procedure that is no longer valid for RHOSP 17 as mistral is not included.
12 February 2024	17.1	Networking	Deploying Ceph in your dynamic routing environment	The first note in the topic has been updated.
5 February 2024	17.1	NFV	Sample DPDK SR-IOV YAML and Jinja2 files	The chapter containing sample YAML files has been updated.
2 February 2024	17.1	NFV	Configuring an SR-IOV deployment	Chapter 7 on SR-IOV has been completely rewritten.
25 January 2024	17.1	Networking	Defining leaf roles and attaching networks	The procedure has changed significantly.
25 January 2024	17.1	NFV	Preventing packet loss by managing RX-TX queue size	The procedure has been rewritten.

Date	17.1 versions impacted	Components	Affected content	Description of change
24 January 2024	17.1	Networking	Migrating the ML2 mechanism driver from OVS to OVN	Now you can migrate to OVN from OVS with VLAN tenant networks and DVR. Also clarified that the environment files shown are just examples, to be replaced with your own files.
23 January 2024	17.1	Security	Managing OpenStack Identity resources	Old and duplicated material is removed from the Identity resources guide
17 January 2024	17.1	Networking	• Deploying the DNS service • Deploying the DNS service with pre-existing BIND 9 servers	A step was added to the two deployment topics instructing administrators to add the name server records (NS records) for the child zones that reside in the DNS server (designate) pool.
15 January 2024	17.1	Edge	• Updating the central location • Deploying edge nodes without storage • Deploying edge sites with hyperconverged storage	A step was added to several procedures to instruct users to re-run the network provisioning command on the central location, if the **network_data.yaml** template includes additional networks which were not included when networks were provisioned for the central location.
17 January 2024	17.1.2	director Operator	Upgrading an overcloud on a Red Hat OpenShift Container Platform cluster with director Operator (16.2 to 17.1)	Added a chapter about how to upgrade an overcloud on RHOCP with director Operator from RHOSP 16.2 to RHOSP 17.1.
17 January 2024	17.1	Upgrades	Known issues that might block an upgrade	Known issues were removed for the following BZs: • BZ#2235621 • BZ#2237743 • BZ#2228818

Date	17.1 versions impacted	Components	Affected content	Description of change
17 January 2024	17.1	Updates	• Validating RHOSP before the undercloud update • Validating RHOSP after the overcloud update	Updated note to describe the **SKIPPED** and **FAILED** statuses that might occur when you run a validation.
17 January 2024	17.1	Upgrades	Overcloud adoption for multi-cell environments	Added a new module that provides an example of adopting the overcloud in multi-cell environments.
12 January 2024	17.1	DCN	Managing separate heat stacks	Removed example of manually creating file that is created automatically
11 January 2024	17.1	Networking	Deploying Ceph in your dynamic routing environment	Added a new topic about how to deploy Red Hat Ceph Storage.
10 January 2024	17.1	Networking	DVR known issues and caveats	Refined an item about DHCP in an ML2/OVS, DVR environment: "For ML2/OVS environments, the DHCP server is not distributed and is deployed on a Controller node. The ML2/OVS neutron DCHP agent, which manages the DHCP server, is deployed in a highly available configuration on the Controller nodes, regardless of the routing design (centralized or DVR)."
10 January 2024	17.1	NFV	Other parameters	Added an IMPORTANT admonition under **VhostuserSocketGroup**.
10 January 2024	17.1	NFV	Tested NICs for NFV	Rewrote topic.
8 January 2024	17.1	Networking	Configuring policy-based routing	Corrected example interface entry: changed **route_table: 2** to **table: 2**.

Date	17.1 versions impacted	Components	Affected content	Description of change
8 January 2024	17.1	NFV	Launching an RT-KVM instance	Corrected a **create flavor** command example.
30 November 2023	17.1	Security	Secure metadef APIs	The concept and procedure have been updated for metadef APIs
30 November 2023	17.1	Networking	Configuring the L2 population driver	The topic, "Configuring the L2 population driver" has been rewritten.
20 November 2023	17.1	Networking	Creating secure HTTP load balancers	Changed prerequisities for procedures in Chapter 9.
15 November 2023	17.1	Security	Implementing TLS-e with Ansible	Added optional step to include **CertmongerKerberosRealm** parameter when the IPA realm does not match the IPA domain.
15 November 2023	17.1	Edge	Deploying distributed compute node architecture with TLS-e	Removed unnecessary (but non-impactful) steps from instructions for TLSe in DCN guide
7 November 2023	17.1	All	Example: Providing feedback on Red Hat documentation	Replaced the Direct Documentation Feedback (DDF) instructions with the **Create Issue** Jira form link. DDF was removed for Red Hat OpenStack Platform, and feedback must now be provided in Jira.
03 November 2023	17.1	Networking	Customizing NIC mappings for pre-provisioned nodes	Modified the topic, "Customizing NIC mappings for pre-provisioned nodes."
03 November 2023	17.1	Networking	Network interface configuration options	Corrected example for a Linux bond.

Date	17.1 versions impacted	Components	Affected content	Description of change
02 November 2023	17.1	Networking	High-level changes in Red Hat OpenStack Platform 17.1	Added item: "In ML2/OVN deployments, you can enable egress minimum and maximum bandwidth policies for hardware offloaded ports."
2 November 2023	17.1	Networking	ML2/OVS to ML2/OVN in-place migration scenarios that have not been validated	Added known issue prohibiting migration to the OVN mechanism driver if your original ML2/OVS environment includes iptables hybrid firewall and trunk ports.
01 November 2023	17.1	Networking	Migration constraints	Re-wrote sub-topic, "Live migration on ML2/OVS deployments."
31 October 2023	17.1	Networking	• Deploying a spine-leaf enabled overcloud • Deploying a spine-leaf enabled overcloud	Removed the VIP definition file, **spine-leaf-vip-data.yaml**, from the **overcloud deploy** command example.
30 October 2023	17.1	Networking	• Adding a composable network • Configuring DNS endpoints	Described how to use the **CloudName{network.name}** definition to set the DNS name for an API endpoint on a composable network that uses a virtual IP.
30 October 2023	17.1	NFV	• Configuring components of OVS hardware offload	Added note about Red Hat Enterprise Linux Traffic Control (TC) subsystem supporting connection tracking (conntrack) helpers or application layer gateways (ALGs).

Date	17.1 versions impacted	Components	Affected content	Description of change
24 October 2023	17.1	Networking	• Configuring Network service availability zones with ML2/OVN • Considerations for networking on DCN architecture	Added an important admonition about router gateway ports.
24 October 2023	17.1	Networking	• Specifying the name that DNS assigns to port • Preparing the undercloud	Added an important admonition about internal DNS resolution for port names.
24 October 2023	17.1	Networking	Adding a new leaf to a spine-leaf deployment	Modified step 8, and added three new steps (9-11).
23 October 2023	17.1	Networking	Performing basic ICMP testing within the ML2/OVN namespace	Clarified login example (step 4).
20 October 2023	17.1	Networking	Exporting the DNS service pool configuration	Updated the procedure to describe how to run the command inside a container.
17 October 2023	17.1	Networking	Setting the subnet for virtual IP addresses	Removed mention of the **VipSubnetMap** parameter, plus some other changes made.
12 October 2023	17.1	Storage	*Creating and managing images*	Content about creating images has been moved to its own chapter called 'Creating RHEL KVM or RHOSP-compatible images'.
05 October 2023	17.1	Networking	*Configuring DNS as a service*	Instances of the tripleo template filename have changed from **enable-designate.yaml** to **designate.yaml**.

Date	17.1 versions impacted	Components	Affected content	Description of change
05 October 2023	17.1	NFV	• Deploying OVN with OVS-DPDK and SR-IOV • Configuring OVS-DPDK parameters • Configuring OVS-DPDK parameters	The step about adding custom resources for OVS-DPDK with the **resource_registry** parameter has been removed.
05 October 2023	17.1	NFV	Tested NICs for NFV	Replaced the filename, **compute-ovs-dpdk.yaml**, with the phrase, "j2 network configuration template."
05 October 2023	17.1	NFV	Configuring NIC partitioning	The YAML file, **os-net-config.yaml**, has been changed to, **roles_data.yaml**.
05 October 2023	17.1	NFV	Registering and enabling repositories	The repository name, **openstack-for-rhel-9-x86_64-rpms**, has been changed to, **openstack-17.1-for-rhel-9-x86_64-rpms**.
04 October 2023	17.1	Networking	• Enabling custom composable networks • Removing an overcloud stack	Replaced networks definition file, **network_data.yaml**, with **network_data_v2.yaml**.
02 October 2023	17.1	Networking	• Configuring the leaf networks • Installing and configuring the undercloud for RHOSP dynamic routing	The **FrrBgpAsn** and **FrrOvnBgpAgentAsn** parameters are now role-based. There is a new parameter, **tripleo_frr_ovn_bgp_agent_enable**.
29 September 2023	17.1	Security	• Enabling FIPS	Corrected procedure so that FIPS images are uploaded to glance.

Date	17.1 versions impacted	Components	Affected content	Description of change
27 September 2023	17.1	Networking	• Constraints for RHOSP dynamic routing	Added item that describes high connectivity downtime during an FRR update for RHOSP dynamic routing environments.
27 September 2023	17.1, 17.0, 16.2, 16.1	Compute	• Adding dynamic metadata to instances	The configuration for dynamic metadata you use in your Compute environment file has been updated.
26 September 2023	17.1	Networking	• Limiting queries to the metadata service	A new procedure, "Limiting queries to the metadata service," has been added to *Configuring Red Hat OpenStack Platform networking*.
25 September 2023	17.1	NFV	• *Configuring network functions virtualization*	*Configuring network functions virtualization* has been updated with content from what was the *Network Functions Virtualization Product Guide*.
22 September 2023	17.1	Deployment	• *Deploying Red Hat OpenStack Platform at scale*	This guide is being reviewed and will be republished on the Customer Portal when the reviewed content is available for enterprise use.
20 September 2023	17.1	Networking	• Creating custom virtual routers with router flavors	Added a chapter for the Technology Preview of the router flavors feature.
19 September 2023	17.1	Updates	• Rebooting Compute nodes	Added note stating that only cold migration is supported when migrating virtual machines from RHEL 9.2 to RHEL 8.4 in a Multi-RHEL environment.

Date	17.1 versions impacted	Components	Affected content	Description of change
19 September 2023	17.1	Upgrades	• Creating roles for Multi-RHEL Compute nodes • Upgrading the Compute node operating system	Updated the parameter that sets the RHEL version on Compute nodes in a Multi-RHEL environment. Removed step to modify the **skip_rhel_release.yaml** file.
18 September 2023	17.1	Edge	• Deploying the central site with storage • Deploying edge nodes without storage • Deploying edge sites with hyperconverged storage • Updating the central location • Delete the DistributedComputeHCI node • Deploying the central controllers without edge storage	Removed mentions of heat template podman.yaml, which is no longer needed.
18 September 2023	17.1	Updates	• Validating RHOSP before the undercloud update • Validating RHOSP after the overcloud update	Added procedures to validate your RHOSP environment before the undercloud update and after the overcloud update.

Date	17.1 versions impacted	Components	Affected content	Description of change
12 September 2023	17.1	Networking	link:https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/17.1/html/configuring_the_compute_service_for_instance_creation/assembly_migrating-virtual-machine-instances-between-compute-nodes_migrating-instances#con_migration-constraints_migrating-instances	Removed the sub-section, "Packet loss on ML2/OVN deployments" from the section, "16.3.2. Migration constraints," in the *Configuring the Compute service for instance creation* guide.
12 September 2023	17.1	NFV	Chapter 6. Preparing network functions virtualization (NFV)	A new chapter, "Chapter 6. Preparing network functions virtualization (NFV)," has been added to the *Framework for upgrades (16.2 to 17.1)* guide.
11 September 2023	17.1	Networking	Chapter 20. Replacing Controller nodes	Changes made to Chapter 20 to address the OVN database partition issue described in BZ 2222543
08 September 2023	17.1	Backup and Restore	Backing up and restoring the undercloud and control plane nodes	*Backing up and restoring the undercloud and control plane nodes* is now published.
07 September 2023	17.1	Networking	QoS rules	To Table 9.1, added footnote (#8) stating that RHOSP does not support QoS for trunk ports.
31 August 2023	17.1	Networking	Configuring the leaf networks and Configuring the leaf networks	Renamed section 4.4 to "Configuring the leaf networks" in both the *Configuring dynamic routing in Red Hat OpenStack Platform* guide and the *Configuring spine-leaf networking* guide.
30 August 2023	17.1	Networking	Overview of allowed address pairs	Added a definition for a virtual port (vport).

Date	17.1 versions impacted	Components	Affected content	Description of change
30 August 2023	17.1	Security	Creating images	Removed deprecated example for building images and replaced with link to image builder documentation
29 August 2023	17.1	Documentation	• *Command line interface reference* • *Configuration reference* • *Overcloud parameters*	New editions published.
28 August 2023	17.1	NFV	Planning for your RT-KVM Compute nodes	The repositories listed in this procedure have changed.
28 August 2023	17.1	NFV	Registering and enabling repositories	The repositories listed in this procedure have changed.
25 August 2023	17.1	Firewall rules for Red Hat OpenStack Platform	Firewall rules for Red Hat OpenStack Platform	New edition published.
22 August 2023	17.1	NFV	Configuring OVS PMD Auto Load Balance	The OVS Poll Mode Driver (PMD) automatic load balancing feature graduated from Technology Preview to full support. Also, the configuration procedure changed.
21 August 2023	17.1	DCN	Considerations for networking on DCN architecture	The RHOSP Load-balancing service (octavia) is no longer listed as unsupported in a DCN environment.
16 August 2023	17.1	Documentation	*Configuring dynamic routing in Red Hat OpenStack Platform*	A new guide for RHOSP 17.1.

Date	17.1 versions impacted	Components	Affected content	Description of change
16 August 2023	17.1	Documentation	• *Network Functions Virtualization Product Guide*	This guide is being reviewed and will be added after the initial release. This list of unpublished guides will be updated when these guides are published.
16 August 2023	17.1	Documentation	• Documentation library updates	Updated the titles for some of the guides from the 17.0 title.
16 August 2023	17.1	Documentation	• *Backing up Block Storage volumes*	This guide has been updated, restructured and rewritten. The cinder CLI commands have been replaced with openstack CLI commands, where possible.

Table 4.3. Documentation library title changes

Previous title	Current title
Bare Metal Provisioning	*Configuring the Bare Metal Provisioning service*
Block Storage Backup Guide	*Backing up Block Storage volumes*
Custom Block Storage Back End Deployment Guide	*Deploying a custom Block Storage back end*
Deployment Recommendations for Specific Red Hat OpenStack Platform Services	Removed.
Installing and managing Red Hat OpenStack Platform with director	*Installing and managing Red Hat OpenStack Platform with director*
Distributed compute node and storage deployment	*Deploying a Distributed Compute Node (DCN) architecture*
External Load Balancing for the Overcloud	Content moved to *Managing high availability services*.
High Availability Deployment and Usage	*Managing high availability services*
High Availability for Compute Instances	*Configuring high availability for instances*
Hyperconverged Infrastructure Guide	*Deploying a hyperconverged infrastructure*

Previous title	Current title
Introduction to the OpenStack Dashboard	*Managing cloud resources with the Openstack Dashboard*
IPv6 networking for the overcloud	*Configuring IPv6 networking for the overcloud*
Keeping Red Hat OpenStack Platform Updated	*Performing a minor update of Red Hat OpenStack Platform*
Network Functions Virtualization Planning and Configuration Guide	*Configuring network functions virtualization*
Network Functions Virtualization Product Guide	Removed. Content will be moved to *Configuring network functions virtualization*.
Networking Guide	*Configuring Red Hat OpenStack Platform networking*
OpenStack Integration Test Suite Guide	*Validating your cloud with the Red Hat OpenStack Platform Integration Test Suite*
Operational Measurements	*Managing overcloud observability*
Partner Integration	Removed.
Product Guide	*Introduction to Red Hat OpenStack Platform*
Recommendations for Large Deployments	*Deploying Red Hat OpenStack Platform at scale*
RHOSP director Operator for OpenShift Container Platform	*Deploying an overcloud in a Red Hat OpenShift Container Platform cluster with director Operator*
Security and Hardening Guide	*Hardening Red Hat OpenStack Platform*
Spine Leaf Networking	*Configuring spine-leaf networking*
Standalone Deployment Guide	Removed.
Storage Guide	*Configuring persistent storage*
Testing Migration of the Networking Service to the ML2/OVN Mechanism Driver	*Migrating to the OVN mechanism driver*
Transitioning to Containerized Services	Removed
Users and Identity Management Guide	*Managing OpenStack Identity resources*
Using Designate for DNS-as-a-Service	*Configuring DNS as a service*

Previous title	Current title
Using Octavia for Load Balancing-as-a-Service	*Configuring load balancing as a service*