



Red Hat OpenStack Platform 17.0

Release Notes

Release details for Red Hat OpenStack Platform 17.0

Red Hat OpenStack Platform 17.0 Release Notes

Release details for Red Hat OpenStack Platform 17.0

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

Legal Notice

Copyright © 2023 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	3
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	4
CHAPTER 1. INTRODUCTION	5
1.1. ABOUT THIS RELEASE	5
1.2. REQUIREMENTS	5
1.3. DEPLOYMENT LIMITS	5
1.4. DATABASE SIZE MANAGEMENT	5
1.5. CERTIFIED DRIVERS AND PLUG-INS	6
1.6. CERTIFIED GUEST OPERATING SYSTEMS	6
1.7. PRODUCT CERTIFICATION CATALOG	6
1.8. BARE METAL PROVISIONING OPERATING SYSTEMS	6
1.9. HYPERVISOR SUPPORT	6
1.10. CONTENT DELIVERY NETWORK (CDN) REPOSITORIES	6
1.10.1. Undercloud repositories	7
1.10.2. Overcloud repositories	8
1.11. PRODUCT SUPPORT	11
1.12. UNSUPPORTED FEATURES	11
CHAPTER 2. TOP NEW FEATURES	13
2.1. BARE METAL SERVICE	13
2.2. BLOCK STORAGE	13
2.3. CEPH STORAGE	14
2.4. COMPUTE	14
2.5. NETWORKING	14
2.6. VALIDATION FRAMEWORK	15
2.7. TECHNOLOGY PREVIEWS	15
CHAPTER 3. RELEASE INFORMATION	17
3.1. RED HAT OPENSTACK PLATFORM 17.0 GA - SEPTEMBER 21, 2022	17
3.1.1. Advisory list	17
3.1.2. Bug Fix	17
3.1.3. Enhancements	19
3.1.4. Technology Preview	21
3.1.5. Release Notes	23
3.1.6. Known Issues	24
3.1.7. Deprecated Functionality	28
3.1.8. Removed Functionality	30
3.2. RED HAT OPENSTACK PLATFORM 17.0.1 MAINTENANCE RELEASE - JANUARY 25, 2023	31
3.2.1. Advisory list	31
3.2.2. Bug Fix	32
3.2.3. Known Issues	33
CHAPTER 4. TECHNICAL NOTES	36
4.1. RHEA-2022:6543 – RELEASE OF COMPONENTS FOR OSP 17.0	36
4.2. RHBA-2023:0271 – RED HAT OPENSTACK PLATFORM 17.0.1 BUG FIX AND ENHANCEMENT ADVISORY	41
CHAPTER 5. DOCUMENTATION CHANGES	45

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.

Using the Direct Documentation Feedback (DDF) function


Use the **Add Feedback** DDF function for direct comments on specific sentences, paragraphs, or code blocks.

1. View the documentation in the *Multi-page HTML* format.
2. Ensure that you see the **Feedback** button in the upper right corner of the document.
3. Highlight the part of text that you want to comment on.
4. Click **Add Feedback**.
5. Complete the **Add Feedback** field with your comments.
6. Optional: Add your email address so that the documentation team can contact you for clarification on your issue.
7. Click **Submit**.

CHAPTER 1. INTRODUCTION

1.1. ABOUT THIS RELEASE

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



NOTE

The Red Hat Enterprise Linux High Availability Add-On is available for Red Hat OpenStack Platform 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 Red Hat OpenStack Platform, see <https://access.redhat.com/site/solutions/509783>.

1.2. REQUIREMENTS

This version of Red Hat OpenStack Platform runs on the most recent fully supported release of Red Hat Enterprise Linux 9.0 Extended Update Support (EUS).

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

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



NOTE

Because Internet Explorer 11 is no longer maintained, expect a degradation of functionality when displaying the dashboard.



NOTE

Before you deploy Red Hat OpenStack Platform, 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, see [Deployment Limits for Red Hat OpenStack Platform](#).

1.4. DATABASE SIZE MANAGEMENT

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

1.5. CERTIFIED DRIVERS AND PLUG-INS

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

1.6. CERTIFIED GUEST OPERATING SYSTEMS

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

1.7. PRODUCT CERTIFICATION CATALOG

For a list of the Red Hat Official Product Certification Catalog, see [Product Certification Catalog](#).

1.8. BARE METAL PROVISIONING OPERATING SYSTEMS

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

1.9. HYPERVISOR SUPPORT

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

This release of the Red Hat OpenStack Platform runs with Bare Metal Provisioning.

Bare Metal Provisioning has been fully supported since the release of Red Hat OpenStack Platform 7 (Kilo). You can use Bare Metal Provisioning to provision bare-metal machines by using common technologies such as PXE and IPMI, to cover a wide range of hardware while supporting pluggable drivers to allow the addition of vendor-specific functionality.

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

1.10. CONTENT DELIVERY NETWORK (CDN) REPOSITORIES

This section describes the repositories required to deploy Red Hat OpenStack Platform 17.0.

You can install Red Hat OpenStack Platform 17.0 through the Content Delivery Network (CDN) using **subscription-manager**.

For more information, see [Planning your undercloud](#).

**WARNING**

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

1.10.1. Undercloud repositories

You run Red Hat OpenStack Platform 17.0 on Red Hat Enterprise Linux 9.0. As a result, you must lock the content from these repositories to the respective Red Hat Enterprise Linux version.

**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).

**NOTE**

Satellite repositories are not listed because RHOSP 17.0 does not support Satellite. Satellite support is planned for a future release. Only Red Hat CDN is supported as a package repository and container registry.

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.

Name	Repository	Description of requirement
Red Hat OpenStack Platform 17.0 for RHEL 9 (RPMs)	openstack-17-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.

1.10.2. Overcloud repositories

You run Red Hat OpenStack Platform 17.0 on Red Hat Enterprise Linux 9.0. As a result, you must lock the content from these repositories to the respective Red Hat Enterprise Linux version.



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).



NOTE

Satellite repositories are not listed because RHOSP 17.0 does not support Satellite. Satellite support is planned for a future release. Only Red Hat CDN is supported as a package repository and container registry. NFV repositories are not listed because RHOSP 17.0 does not support NFV.

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.

Name	Repository	Description of requirement
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 17 for RHEL 9 x86_64 (RPMs)	openstack-17-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 5 for RHEL 9 x86_64 (RPMs)	rhceph-5-tools-for-rhel-9-x86_64-rpms	Tools for Red Hat Ceph Storage 5 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 17 for RHEL 9 x86_64 (RPMs)	openstack-17-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 5 for RHEL 9 x86_64 (RPMs)	rhceph-5-tools-for-rhel-9-x86_64-rpms	Tools for Red Hat Ceph Storage 5 for Red Hat Enterprise Linux 9.

Real Time Compute repositories

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

Name	Repository	Description of requirement
Red Hat Enterprise Linux 9 for x86_64 - Real Time (RPMs)	rhel-9-for-x86_64-rt-rpms	Repository for Real Time KVM (RT-KVM). Contains packages to enable the real time kernel. Enable this repository for all Compute nodes targeted for RT-KVM. NOTE: You need a separate subscription to a Red Hat OpenStack Platform for Real Time SKU to access this repository.

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 17 Director Deployment Tools for RHEL 9 x86_64 (RPMs)	openstack-17-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-for-rhel-9-x86_64-rpms repository.
Red Hat OpenStack Platform 17 for RHEL 9 x86_64 (RPMs)	openstack-17-for-rhel-9-x86_64-rpms	Packages to help director configure Ceph Storage nodes. This repository is included with combined OpenStack Platform and Ceph Storage subscriptions. If you use a standalone Ceph Storage subscription, use the openstack-17-deployment-tools-for-rhel-9-x86_64-rpms repository.
Red Hat Ceph Storage Tools 5 for RHEL 9 x86_64 (RPMs)	rhceph-5-tools-for-rhel-9-x86_64-rpms	Provides tools for nodes to communicate with the Ceph Storage cluster.

Name	Repository	Description of requirement
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.

1.11. 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>.

1.12. UNSUPPORTED FEATURES

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

- There is no NFV support in RHOSP 17.0.
- Custom policies, which includes modification of **policy.json** files either manually or through any **Policies** heat parameters. Do not modify the default policies unless the documentation contains explicit instructions to do so.
- Containers are not available for the following packages, therefore they are not supported in RHOSP:
 - **nova-serialproxy**
 - **nova-spicehtml5proxy**

- File injection of personality files to inject user data into virtual machine instances. Instead, cloud users can pass data to their instances by using the **--user-data** option to run a script during instance boot, or set instance metadata by using the **--property** option when launching an instance. For more information, see [Creating a customized instance](#).
- 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](#)
 - [Intel® Product Change Notification #119311-00](#)

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.

2.1. BARE METAL SERVICE

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

Provision hardware before deploying the overcloud

In Red Hat OpenStack Platform 17.0, you must provision the bare metal nodes and the physical networks resources for the overcloud before deploying the overcloud. The **openstack overcloud deploy** command no longer provisions the hardware. For more information, see [Provisioning and deploying your overcloud](#).

New network definition file format

In Red Hat OpenStack Platform 17.0, you configure your network definition files by using ansible jinja2 templates instead of heat templates. For more information, see [Configuring overcloud networking](#).

Whole disk images are the default overcloud image

The default **overcloud-full** flat partition images have been updated to **overcloud-hardened-uefi-full** whole disk images. The whole disk image is a single compressed qcow2 image that contains the following elements:

- A partition layout containing UEFI boot, legacy boot, and a root partition.
- The root partition contains a single lvm group with logical volumes of different sizes that are mounted at `/`, `/tmp`, `/var`, `/var/log`, and so on.
When you deploy a whole-disk image, `ironic-python-agent` copies the whole image to the disk without any bootloader or partition changes.

UEFI Boot by default

The default boot mode of bare metal nodes is now UEFI boot, because the Legacy BIOS boot feature is unavailable in new hardware.

2.2. BLOCK STORAGE

This section outlines the top new features for the Block Storage (cinder) service.

Support for automating multipath deployments

You can specify the location of your multipath configuration file for your overcloud deployment.

Project-specific default volume types

For complex deployments, project administrators can define a default volume type for each project (tenant).

If you create a volume and do not specify a volume type, then Block Storage uses the default volume type. You can use the Block Storage (cinder) configuration file to define the general default volume type that applies to all your projects (tenants). But if your deployment uses project-specific volume types, ensure that you define default volume types for each project. In this case, Block Storage uses the project-specific volume type instead of the general default volume type. For more information, see [Defining a project-specific default volume type](#).

2.3. CEPH STORAGE

This section outlines the top new features for Ceph Storage.

Greater security for Ceph client Shared Files Systems service (manila) permissions

The Shared File Systems service CephFS drivers (native CephFS and CephFS through NFS) now interact with Ceph clusters through the Ceph Manager **Volumes** interface. The Ceph client user configured for the Shared Files Systems service no longer needs to be as permissive. This feature makes Ceph client user permissions for the Shared Files Systems service more secure.

Ceph Object Gateway (RGW) replaces Object Storage service (swift)

When you use Red Hat OpenStack Platform (RHOSP) director to deploy Ceph, director enables Ceph Object Gateway (RGW) object storage, which replaces the Object Storage service (swift). All other services that normally use the Object Storage service can start using RGW instead without additional configuration.

Red Hat Ceph Storage cluster deployment in new environments

In new environments, the Red Hat Ceph Storage cluster is deployed first, before the overcloud, using director and the `openstack overcloud ceph deploy` command. You now use `cephadm` to deploy Ceph, because deployment with `ceph-ansible` is deprecated. For more information about deploying Ceph, see [Deploying Red Hat Ceph Storage and Red Hat OpenStack Platform together with director](#). This document replaces *Deploying an overcloud with containerized Red Hat Ceph*. A Red Hat Ceph Storage cluster that you deployed without RHOSP director is also supported. For more information, see [Integrating an Overcloud with an Existing Red Hat Ceph Storage Cluster](#).

Support for creating shares from snapshots

You can create a new share from a snapshot to restore snapshots by using the Shared File Systems service (manila) CephFS back ends: native CephFS and CephFS through NFS.

2.4. COMPUTE

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

Support for attaching and detaching SR-IOV devices to an instance

Cloud users can create a port that has an SR-IOV vNIC, and attach the port to an instance when there is a free SR-IOV device on the host on the appropriate physical network, and the instance has a free PCIe slot. For more information, see [Attaching a port to an instance](#).

Support for creating an instance with NUMA-affinity on the port

Cloud users can create a port that has a NUMA affinity policy, and attach the port to an instance. For more information, see [Creating an instance with NUMA affinity on the port](#).

Q35 is the default machine type

The default machine type for each host architecture is Q35 (**pc-q35-rhel9.0.0**) for new Red Hat OpenStack Platform 17.0 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.

2.5. NETWORKING

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

Active/Active clustered database service model improves OVS database read performance and fault tolerance

Starting in RHOSP 17.0, RHOSP ML2/OVN deployments use a clustered database service model that applies the Raft consensus algorithm to enhance performance of OVS database protocol traffic and provide faster, more reliable failover handling. The clustered database service model replaces the pacemaker-based, active/backup model.

A clustered database operates on a cluster of at least three database servers on different hosts. Servers use the Raft consensus algorithm to synchronize writes and share network traffic continuously across the cluster. The cluster elects one server as the leader. All servers in the cluster can handle database read operations, which mitigates potential bottlenecks on the control plane. Write operations are handled by the cluster leader.

If a server fails, a new cluster leader is elected and the traffic is redistributed among the remaining operational servers. The clustered database service model handles failovers more efficiently than the pacemaker-based model did. This mitigates related downtime and complications that can occur with longer failover times.

The leader election process requires a majority, so the fault tolerance capacity is limited by the highest odd number in the cluster. For example, a three-server cluster continues to operate if one server fails. A five-server cluster tolerates up to two failures. Increasing the number of servers to an even number does not increase fault tolerance. For example, a four-server cluster cannot tolerate more failures than a three-server cluster.

Most RHOSP deployments use three servers.

Clusters larger than five servers also work, with every two added servers allowing the cluster to tolerate an additional failure, but write performance decreases.

The clustered database model is the default in RHOSP 17.0 deployments. You do not need to perform any configuration steps.

Designate DNSaaS

In Red Hat OpenStack Platform (RHOSP) 17.0, the DNS service (designate) is now fully supported. Designate is an official OpenStack project that provides DNS-as-a-Service (DNSaaS) implementation and enables you to manage DNS records and zones in the cloud. The DNS service provides a REST API, and is integrated with the RHOSP Identity service (keystone) for user management. Using RHOSP director you can deploy BIND instances to contain DNS records, or you can integrate the DNS service into an existing BIND infrastructure. (Integration with an existing BIND infrastructure is a technical preview feature.) In addition, director can configure DNS service integration with the RHOSP Networking service (neutron) to automatically create records for virtual machine instances, network ports, and floating IPs. For more information, see [Using Designate for DNS-as-a-Service](#).

2.6. VALIDATION FRAMEWORK


This section outlines the top new features for the Validation Framework.

User-created validations through the CLI

In Red Hat OpenStack Platform (RHOSP) 17.0, you can create your own personalized validation with the validation init command. Execution of the command results in a template for a new validation. You can edit the new validation role to suit your requirements.

2.7. TECHNOLOGY PREVIEWS

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

**NOTE**

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

Border Gateway Protocol (BGP)

In Red Hat OpenStack Platform (RHOSP) 17.0, a technology preview is available for Border Gateway Protocol (BGP) to route the control plane, floating IPs, and workloads in provider networks. By using BGP advertisements, you do not need to configure static routes in the fabric, and RHOSP can be deployed in a pure Layer 3 data center. RHOSP uses Free Range Routing (FRR) as the dynamic routing solution to advertise and withdraw routes to control plane endpoints as well as to VMs in provider networks and Floating IPs.

Integrating existing BIND servers with the DNS service

In Red Hat OpenStack Platform (RHOSP) 17.0, a technology preview is available for integrating the RHOSP DNS service (designate) with an existing BIND infrastructure. For more information, see [Configuring existing BIND servers for the DNS service](#).

CHAPTER 3. RELEASE INFORMATION

These release notes highlight technology preview items, recommended practices, known issues, and deprecated functionality that you should consider when you deploy this release of Red Hat OpenStack Platform.

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

3.1. RED HAT OPENSTACK PLATFORM 17.0 GA - SEPTEMBER 21, 2022

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

3.1.1. Advisory list

This release includes the following advisories:

[RHEA-2022:6543](#)

Release of components for Red Hat OpenStack Platform 17.0 (Wallaby)

[RHEA-2022:6544](#)

Release of containers for Red Hat OpenStack Platform 17.0 (Wallaby)

[RHEA-2022:6545](#)

Red Hat OpenStack Platform 17.0 RHEL 9 deployment images (qcow2 tarballs)

[RHEA-2022:6546](#)

Red Hat OpenStack Platform 17.0 (Wallaby) RHEL 9 deployment images (RPMs)

3.1.2. Bug Fix

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

[BZ#1374002](#)

Before this update, a misconfiguration of communication parameters between the DNS service (designate) worker and deployed BIND instances caused Red Hat OpenStack Platform (RHOSP) 17.0 Beta deployments that have more than one Controller node to fail. With this update, this issue has been resolved, and you can now use the DNS service in a deployment with more than one Controller node.

[BZ#1801931](#)

Before this update, the help text for the **max_disk_devices_to_attach** parameter did not state that **0** is an invalid value. Also, when the **max_disk_devices_to_attach** parameter was set to **0**, the **nova-compute** service started when it should have failed. With this update, the **max_disk_devices_to_attach** parameter help option text states that a value of **0** is invalid, and if **max_disk_devices_to_attach** is set to **0**, the **nova-compute** service will now log an error and fail to start.

[BZ#1883326](#)

Before this update, an issue existed with PowerFlex storage-assisted volume migration when volume migration was performed without conversion of volume type in cases where it should have been converted to thin from thick provisioned. With this update, this issue is fixed.

[BZ#1888069](#)

Before this update, Supermicro servers in UEFI mode would reboot from the network instead of from the local hard disk, causing a failed boot. With this update, Ironic sends the correct raw IPMI commands that request UEFI "boot from hard disk." Booting Supermicro nodes in UEFI mode with IPMI now works as expected.

BZ#1944586

This update fixes a bug that incorrectly redirected registered non-stdout callback output from various Ansible processes to the validations logging directory. Output of other processes is no longer stored in validations logging directory. VF callbacks no longer receive information about plays, unless requested.

BZ#1984556

The collectd smart plugin requires the CAP_SYS_RAWIO capability. CAP_SYS_RAWIO is not present by default in the configuration, and before this update, you could not add it. With this update, you can use the CollectdContainerAdditionalCapAdd parameter to add CAP_SYS_RAWIO. Enter the following parameter value assignment in an environment file.

Example

```
parameter_defaults:
  CollectdExtraPlugins:
    - smart

  CollectdContainerAdditionalCapAdd: "CAP_SYS_RAWIO"
```

BZ#1991657

Before this update, baremetal node introspection failed with an error and did not retry, when the node had a transient lock on it. With this update, you can perform introspection even when the node has a lock.

BZ#2050773

Before this update, if an operator defined a custom value for the **volume:accept_transfer** policy that referred to the project_id of the user making the volume transfer accept request, the request would fail. This update removes a duplicate policy check that incorrectly compared the project_id of the requestor to the project_id associated with the volume before transfer. The check done at the Block Storage API layer will now function as expected.

BZ#2064019

Before this update, network interruptions caused a bare metal node's power state to become **None**, and enter the **maintenance** state. This is due to Ironic's connection cache of Redfish node sessions entering a stale state and not being retried. This state cannot be recovered without restarting the Ironic service. With this update, the underlying REST client has been enhanced to return specific error messages. These error messages are used by Ironic to invalidate cached sessions.

BZ#2101937

With this fix, traffic is distributed on VLAN provider networks in ML2/OVN deployments. Previously, traffic on VLAN provider networks was centralized even with the Distributed Virtual Router (DVR) feature enabled.

BZ#2121098

Before this update in Red Hat OpenStack Platform (RHOSP) 17.0 Beta, Networking service (neutron) requests could fail with a **504 Gateway Time-out** if they occurred when the Networking service reconnected to **ovsdb-server**. These reconnections could happen during failovers or through **ovsdb-server** leader transfers during database compaction.

If neutron debugging was enabled, the Networking service rapidly logged a large number of OVSDDB transaction returned TRY_AGAIN" DEBUG messages, until the transaction timed out with an exception.

With this update, the reconnection behavior is fixed to handle this condition, with a single retry of the transaction until a successful reconnection.

3.1.3. Enhancements

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

BZ#1689706

This enhancement includes OpenStack CLI (OSC) support for Block Storage service (cinder) API 3.42. This allows OSC to extend an online volume.

BZ#1699454

With this update, you can restore snapshots with the CephFS Native and CephFS with NFS backends of the Shared File Systems service (manila) by creating a new share from a snapshot.

BZ#1752776

In Red Hat OpenStack Platform (RHOSP) 17.0 GA, non-admin users have access to new parameters when they run the **openstack server list** command:

- --availability-zone <az_name>
- --config-drive
- --key-name <key_name>
- --power-state <state>
- --task-state <state>
- --vm-state <state>
- --progress <percent_value>
- --user <name_or_ID>

For more information, see [server list](#).

BZ#1758161

With this update, Red Hat OpenStack Platform director deployed Ceph includes the RGW daemon, replacing the Object Storage service (swift) for object storage. To keep the Object Storage service, use the **cephadm-rbd-only.yaml** file instead of **cephadm.yaml**.

BZ#1813560

With this update, the Red Hat OpenStack Platform (RHOSP) 17 Octavia amphora image now includes HAProxy 2.4.x as distributed in Red Hat Enterprise Linux (RHEL) 9. This improves the performance of Octavia load balancers; including load balancers using flavors with more than one vCPU core.

BZ#1839169

With this update, **cephadm** and **orchestrator** replace ceph-ansible. You can use director with cephadm to deploy the ceph cluster and additional daemons, and use a new `tripleo-ansible` role to configure and enable the Ceph backend.

BZ#1848153

With this update, you can now use Red Hat OpenStack Platform director to configure the etcd service to use TLS endpoints when deploying TLS-everywhere.

BZ#1903610

This enhancement adds the MemcachedMaxConnections parameter. You can use MemcachedMaxConnections to control the maximum number of memcache connections.

BZ#1904086

With this enhancement, you can view a volume Encryption Key ID using the cinder client command 'cinder --os-volume-api-version 3.64 volume show <volume_name>'. You must specify microversion 3.64 to view the value.

BZ#1944872

This enhancement adds the '--limit' argument to the 'openstack tripleo validator show history' command. You can use this argument to show only a specified number of the most recent validations.

BZ#1946956

This enhancement changes the default machine type for each host architecture to Q35 (**pc-q35-rhel9.0.0**) for new Red Hat OpenStack Platform 17.0 deployments. The Q35 machine type provides several benefits and improvements, including live migration of instances between different RHEL 9.x minor releases, and the native PCIe hotplug that is faster than the ACPI hotplug used by the **i440fx** machine type.

BZ#1946978

With this update, the default machine type is RHEL9.0-based Q35 **pc-q35-rhel9.0.0**, with the following enhancements:

- Live migration across RHEL minor releases.
- Native PCIe hotplug. This is also ACPI-based like the previous i440fx machine type.
- Intel input-output memory management unit (IOMMU) emulation helps protect guest memory from untrusted devices that are directly assigned to the guest.
- Faster SATA emulation.
- Secure boot.

BZ#1954103

With this enhancement you can use the PluginInstanceFormat parameter for collectd to specify more than one value.

BZ#1954274

This enhancement improves the operating performance of the Bare Metal Provisioning service (ironic) to optimize the performance of large workloads.

BZ#1959707

In Red Hat OpenStack Platform (RHOSP) 17.0 GA, the **openstack tripleo validator show** command has a new parameter, **--limit <number>**, that enables you to limit the number of validations that TripleO displays. The default value is to display the last 15 validations.

For more information, see [tripleo validator show history](#).

BZ#1971607

With this update, the Validation Framework provides a configuration file in which you can set parameters for particular use. You can find an example of this file at the root of the code source or in the default location: **/etc/validation.cfg**.

You can use the default file in `/etc/` or use your own file and provide it to the CLI with the argument `-config`.

When you use a configuration file there is an order for the variables precedence. The following order is the order of variable precedence:

- User's CLI arguments
- Configuration file
- Default interval values

BZ#1973356

This security enhancement reduces the user privilege level required by the OpenStack Shared File System service (manila). You no longer need permissions to create and manipulate Ceph users, because the Shared File Systems service now uses the APIs exposed by the **Ceph Manager** service for this purpose.

BZ#2041429

You can now pre-provision bare metal nodes in your application by using the **overcloud node [un]provision** command.

3.1.4. Technology Preview

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

BZ#1884782

In Red Hat OpenStack Platform (RHOSP) 17.0 GA, a technology preview is available for integration between the RHOSP Networking service (neutron) ML2/OVN and the RHOSP DNS service (designate). As a result, the DNS service does not automatically add DNS entries for newly created VMs.

BZ#1896551

In Red Hat OpenStack Platform (RHOSP) 17.0, a technology preview is available for Border Gateway Protocol (BGP) to route the control plane, floating IPs, and workloads in provider networks. By using BGP advertisements, you do not need to configure static routes in the fabric, and RHOSP can be deployed in a pure Layer 3 data center. RHOSP uses Free Range Routing (FRR) as the dynamic routing solution to advertise and withdraw routes to control plane endpoints as well as to VMs in provider networks and Floating IPs.

BZ#1901686

In Red Hat OpenStack Platform 17.0, secure role-based access control (RBAC) is available for the Load-balancing service (octavia) as a technology preview.

BZ#1901687

In Red Hat OpenStack Platform 17.0, Secure RBAC is available for the DNS service (designate) as a technology preview.

BZ#2008274

In Red Hat OpenStack Platform 17.0, a technology preview is available for integrating the DNS service (designate) with a pre-existing DNS infrastructure that uses BIND 9. For more information, see [Deploying the DNS service with pre-existing BIND 9 servers](#)

BZ#2120392

In Red Hat OpenStack Platform 17.0, a technology preview is available for creating single NUMA node instances that have both pinned and floating CPUs.

BZ#2120407

In Red Hat OpenStack Platform 17.0, a technology preview is available for live migrating, unshelving and evacuating an instance that uses a port that has resource requests, such as a guaranteed minimum bandwidth QoS policy.

BZ#2120410

In Red Hat OpenStack Platform 17.0, a technology preview is available for Compute service scheduling based on routed networks. Network segments are reported to the Placement service as host aggregates. The Compute service includes the network segment information in the Placement service query to ensure that the selected host is connected to the correct network segment. This feature enables more accurate scheduling through better tracking of IP availability and locality, and more accurate instance migration, resizing, or unshelving through awareness of the routed network IP subnets.

BZ#2120743

In Red Hat OpenStack Platform 17.0, a technology preview is available for rescuing an instance booted from a volume.

BZ#2120746

In Red Hat OpenStack Platform 17.0, a technology preview is available to define custom inventories and traits in a declarative **provider.yaml** configuration file. Cloud operators can model the availability of physical host features by using custom traits, such as **CUSTOM_DIESEL_BACKUP_POWER**, **CUSTOM_FIPS_COMPLIANT**, and **CUSTOM_HPC_OPTIMIZED**. They can also model the availability of consumable resources by using resource class inventories, such as **CUSTOM_DISK_IOPS**, and **CUSTOM_POWER_WATTS**. Cloud operators can use the ability to report specific host information to define custom flavors that optimize instance scheduling, particularly when used in collaboration with reserving hosts by using isolated aggregates. Defining a custom inventory prevents oversubscription of Power IOPS and other custom resources that an instance consumes.

BZ#2120756

In Red Hat OpenStack Platform 17.0, a technology preview is available to configure counting of 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#2120757

In Red Hat OpenStack Platform 17.0, a technology preview is available for requesting that images are pre-cached on Compute nodes in a host aggregate, when using microversion 2.81. To reduce boot time, you can request that a group of hosts within an aggregate fetch and cache a list of images.

BZ#2120761

In Red Hat OpenStack Platform 17.0, a technology preview is available to use traits and the Placement service to prefilter hosts by using the supported device model traits declared by the virt drivers.

BZ#2128042

In Red Hat OpenStack Platform 17.0, a technology preview is available for Compute node support of multiple NVIDIA vGPU types for each physical GPU.

BZ#2128056

In Red Hat OpenStack Platform 17.0, a technology preview is available for cold migrating and resizing instances that have vGPUs.

For a known issue affecting the vGPU Technology Preview, see https://bugzilla.redhat.com/show_bug.cgi?id=2116979.

BZ#2128070

In Red Hat OpenStack Platform 17.0, a technology preview is available for creating an instance with a VirtIO data path acceleration (VDPA) interface.

3.1.5. Release Notes

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

BZ#1767084

With this update, the CephFS drivers in the OpenStack Shared File Systems service (manila) are updated so that you can manage provisioning and storage lifecycle operations by using the Ceph Manager API. When you create new file shares, the shares are created in a new format that is quicker for creating, deleting and operations. This transition does not affect pre-existing file shares.

BZ#1813573

This enhancement includes Octavia support for object tags. This allows users to add metadata to load balancer resources and filter query results based on tags.

BZ#2013120

With this update, you can supply a new argument **--skiplist** to the **validation run** command. Use this command with a **yaml** file containing services to skip when running validations.

BZ#2090813

The data collection service (Ceilometer) is supported for collection of Red Hat OpenStack Platform (RHOSP) telemetry and events. Ceilometer is also supported for the transport of those data points to the metrics storage service (gnocchi) for the purposes of autoscaling, and delivery of metrics and events to Service Telemetry Framework (STF) for RHOSP monitoring.

BZ#2111015

In an ML2/OVS deployment, Open vSwitch (OVS) does not support offloading OpenFlow rules that have the **skb_priority**, **skb_mark**, or output queue fields set. Those fields are needed to provide quality-of-service (QoS) support for virtio ports.

If you set a minimum bandwidth rule for a virtio port, the Neutron Open vSwitch agent marks the traffic of this port with a Packet Mark Field. As a result, this traffic cannot be offloaded, and it affects the traffic in other ports. If you set a bandwidth limit rule, all traffic is marked with the default 0 queue, which means no traffic can be offloaded.

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

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

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

BZ#2111527

In RHOSP 17.0 you must use Ceph containers based on RHCSv5.2 GA content.

BZ#2117229

Previously, the **collectd** processes plugin was enabled by default, without a list of processes to watch. This would cause messages in collectd logs like "procs_running not found". With this update, the **collectd** processes plugin is removed from the list of collectd plugins that are installed and enabled by default. You can enable the plugin by adding it to the configuration.

3.1.6. Known Issues

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

BZ#2126476

NFV is not supported in RHOSP 17.0. Do not deploy NFV use cases in RHOSP 17.0.

BZ#1966157

There is a limitation when using ML2/OVN with **provider:network_type geneve** with a Mellanox adapter on a Compute node that has more than one instance on the geneve network. The floating IP of only one of the instances will be reachable. You can track the progress of the resolution on this Bugzilla ticket.

BZ#2085583

There is currently a known issue wherein long-running operations can cause the **ovsdb** connection to time out causing reconnects. These time outs can then cause the **nova-compute** agent to become unresponsive. Workaround: You can use the command-line client instead of the default native python bindings. Use the following parameters in your heat templates to use the command-line client:

```
parameter_defaults:
  ComputeExtraConfig:
    nova:os_vif_ovs:ovsdb_interface => 'vsctl'
```

BZ#2091076

Before this update, the health check status script failed because it relied on the podman log content that was no longer available. Now the health check script uses the podman socket instead of the podman log.

BZ#2105291

There is currently a known issue where 'undercloud-heat-purge-deleted' validation fails. This is because it is not compatible with Red Hat OpenStack Platform 17. Workaround: Skip 'undercloud-heat-purge-deleted' with '--skip-list' to skip this validation.

BZ#2104979

A known issue in RHOSP 17.0 prevents the default mechanism for selecting the hypervisor fully qualified domain name (FQDN) from being set properly if the **resource_provider_hypervisors** heat parameter is not set. This causes the SRIOV or OVS agent to fail to start.

Workaround: Specify the hypervisor FQDN explicitly in the heat template. The following is an example of setting this parameter for the SRIOV agent:

ExtraConfig: neutron::agents::ml2::sriov::resource_provider_hypervisors: "enp7s0f3:%{hiera("fqdn_canonical")},enp5s0f0:%{hiera("fqdn_canonical")}".

BZ#2107896

There is currently a known issue that causes tuned kernel configurations to not be applied after initial provisioning.

Workaround: You can use the following custom playbook to ensure that the tuned kernel command line arguments are applied. Save the following playbook as **/usr/share/ansible/tripleo-playbooks/cli-overcloud-node-reset-blscfg.yaml** on the undercloud node:

```
- name: Reset BLSCFG of compute node(s) meant for NFV deployments
  hosts: allovercloud
  any_errors_fatal: true
  gather_facts: true

  pre_tasks:
    - name: Wait for provisioned nodes to boot
      wait_for_connection:
        timeout: 600
        delay: 10

  tasks:
    - name: Reset BLSCFG flag in grub file, if it is enabled
      become: true
      lineinfile:
        path: /etc/default/grub
        line: "GRUB_ENABLE_BLSCFG=false"
        regexp: "^GRUB_ENABLE_BLSCFG=.*"
        insertafter: '^GRUB_DISABLE_RECOVERY.*'
```

Configure the role in the node definition file, **overcloud-baremetal-deploy.yaml**, to run the **cli-overcloud-node-reset-blscfg.yaml** playbook before the playbook that sets the **kernelargs**:

```
- name: ComputeOvsDpdkSriov
  count: 2
  hostname_format: computeovsdpdkrsriov-%index%
  defaults:
    networks:
      - network: internal_api
        subnet: internal_api_subnet
      - network: tenant
        subnet: tenant_subnet
      - network: storage
        subnet: storage_subnet
    network_config:
      template: /home/stack/osp17_ref/nic-configs/computeovsdpdkrsriov.j2
    config_drive:
      cloud_config:
        ssh_pwauth: true
        disable_root: false
      chpasswd:
        list: |-
          root:12345678
        expire: False
  ansible_playbooks:
```

```

- playbook: /usr/share/ansible/tripleo-playbooks/cli-overcloud-node-reset-blscfg.yaml
- playbook: /usr/share/ansible/tripleo-playbooks/cli-overcloud-node-kernelargs.yaml
  extra_vars:
    reboot_wait_timeout: 600
    kernel_args: 'default_hugepagesz=1GB hugepagesz=1G hugepages=32 iommu=pt
intel_iommu=on isolcpus=1-11,13-23'
    tuned_profile: 'cpu-partitioning'
    tuned_isolated_cores: '1-11,13-23'
- playbook: /usr/share/ansible/tripleo-playbooks/cli-overcloud-openvswitch-dpdk.yaml
  extra_vars:
    memory_channels: '4'
    lcore: '0,12'
    pmd: '1,13,2,14,3,15'
    socket_mem: '4096'
    disable_emc: false
    enable_tso: false
    revalidator: ""
    handler: ""
    pmd_auto_lb: false
    pmd_load_threshold: ""
    pmd_improvement_threshold: ""
    pmd_rebal_interval: ""
    nova_postcopy: true

```

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#2112988

There is currently a known issue where the Swift API does not work and returns a 401 error when multiple Controller nodes are deployed and Ceph is enabled.

A workaround is available at <https://access.redhat.com/solutions/6970061>.

BZ#2116529

Live migration fails when executing the QEMU command **migrate-set-capabilities**. This is because the post-copy feature that is enabled by default is not supported.

Choose one of the following workaround options:

- Workaround Option 1: Set **vm.unprivileged_userfaultfd = 1** on Compute nodes to enable post-copy on the containerized libvirt:
 - Make a new file: **\$ touch /etc/sysctl.d/50-userfault.conf.**
 - Add **vm.unprivileged_userfaultfd = 1** to **/etc/sysctl.d/50-userfault.conf.**
 - Load the file: **\$ sysctl -p /etc/sysctl.d/50-userfault.conf.**
- Workaround Option 2: Set the **sysctl** flag through director, by setting the **ExtraSysctlSettings** parameter.
- Workaround Option 3: Disable the post-copy feature completely, by setting the **NovaLiveMigrationPermitPostCopy** parameter to **false**.

BZ#2116979

When using the Technology Preview vGPU support features, a known issue prevents **mdev** devices from being freed when stopping, moving, or deleting vGPU instances in RHOSP 17. Eventually, all **mdev** devices become consumed, and additional instances with vGPUs cannot be created on the compute host.

BZ#2116980

If you launch a vGPU instance in RHOSP 17 you cannot delete it, stop it, or move it. When an instance with a vGPU is deleted, migrated off its compute host, or stopped, the vGPU's underlying mdev device is not cleaned up. If this happens to enough instances, all available mdev devices will be consumed, and no further instances with vGPUs can be created on that compute host.

BZ#2120383

There is currently a known issue when creating instances that have an emulated Trusted Platform Module (TPM) device. Workaround: Disable Security-Enhanced Linux (SELinux).

BZ#2120398

There is currently a known issue with deploying multi-cell and multi-stack overclouds on RHOSP 17. This is a regression with no workaround, therefore the multi-cell and multi-stack overcloud features are not available in RHOSP 17.0.

BZ#2120766

There is currently a known issue with the RHEL firmware definition file missing from some machine types, which causes the booting of instances with an image firmware of UEFI to fail with a `UEFIUnsupported` exception. This issue is being addressed by https://bugzilla.redhat.com/show_bug.cgi?id=2109644. There is also a known issue when **mem_encryption=on** in the kernel args of an AMD SEV Compute node, that results in the Compute node kernel hanging after a reboot and not restarting. There is no workaround for these issues, therefore the AMD SEV feature is not available in RHOSP 17.0.

BZ#2120773

There is currently a known issue with shutting down and restarting instances after a Compute node reboot on RHOSP 17. When a Compute node is rebooted, the automated process for gracefully shutting down the instance fails, which causes the instance to have less time to shut down before the system forces them to stop. The results of the forced stop may vary. Ensure you have fresh backups for all critical workloads before rebooting Compute nodes.

BZ#2121752

Because of a performance issue with the new socket NUMA affinity policy for PCI passthrough devices and SR-IOV interfaces, the **socket** NUMA affinity policy is not supported in RHOSP 17.0.

BZ#2124294

Sensubility does not have permission to access `/run/podman/podman.sock`, which causes the container health check to fail to send the service container status data to Service Telemetry Framework (STF).

Workaround: Run the following command on all overcloud nodes after deployment: **sudo podman exec -it collectd setfacl -R -m u:collectd:rwx /run/podman**

Result: User collectd gets access to `/run/podman` path recursively allowing sensubility to connect to podman.

BZ#2125159

In Red Hat OpenStack Platform (RHOSP) 17.0 GA, there is a known issue where ML2/OVN deployments fail to automatically create DNS records with the RHOSP DNS service (designate). The cause for this problem is that the required Networking service (neutron) extension, **dns_domain_ports**, is not present.

Workaround: currently there is no workaround, but the fix has been targeted for a future RHOSP release.

BZ#2126810

In Red Hat OpenStack Platform (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#2127965

In Red Hat OpenStack Platform (RHOSP) 17.0 GA, there is a known issue where the Free Range Router (FRR) container does not start after the host on which it resides is rebooted. This issue is caused by a missing file in the BGP configuration.

Workaround: create the file, **/etc/tmpfiles.d/run-frr.conf**, and add the following line:

```
d /run/frr 0750 root root - -
```

After you make this change, **tmpfiles** recreates **/run/frr** after each reboot and the FRR container can start.

BZ#2128928

Integration with Red Hat Satellite is not supported in RHOSP 17.0. Only Red Hat CDN is supported as a package repository and container registry. Satellite support will resume in a future release.

BZ#2120377

You cannot use the UEFI Secure Boot feature because there is currently a known issue with UEFI boot for instances. This is due to an underlying RHEL issue.

BZ#2120384

You cannot create Windows Server 2022 instances on RHOSP because they require vTPM support, which is not currently available.

BZ#2152218

There is currently a known issue when attaching a volume to an instance, or detaching a volume from an instance, when the instance is in the process of booting up or shutting down. You must wait until the instance is fully operational, or fully stopped, before attaching or detaching a volume.

BZ#2153815

There is currently a known issue with creating instances when the instance flavor includes resource usage extra specs, **quota:cpu_***. On RHOSP 17.0, attempts to create an instance with a flavor that limits the CPU quotas encounter the following error: "Requested CPU control policy not supported by host". This error is raised on RHOSP 17.0 on RHEL 9 because the Compute service assumes that the host is running **cggroups** instead of **cggroups-v2**, therefore it incorrectly detects that the host does not support resource usage extra specs.

BZ#2162242

There is currently a known issue with CPU pinning on RHEL 9 kernels older than **kernel-5.14.0-70.43.1.el9_0** that causes soft and hard CPU affinity on all existing **cggroups** to be reset when a new **cgroup** is created. This issue is being addressed in https://bugzilla.redhat.com/show_bug.cgi?id=2143767. To use CPU pinning, update your kernel to **kernel-5.14.0-70.43.1.el9_0** or newer and reboot the host.

3.1.7. Deprecated Functionality

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

BZ#1874778

In Red Hat OpenStack Platform 17.0, the **iscsi** deployment interface has been deprecated. The default deployment interface is now **direct**. Bug fixes and support are provided while the feature is deprecated but Red Hat will not implement new feature enhancements. In a future release, the interface will be removed.

BZ#1946898

In Red Hat OpenStack Platform 17.0, the QEMU **i440fx** machine type has been deprecated. The default machine type is now Q35, **pc-q35-rhel9.0.0**. While the **pc-i440fx-*** machine types are still available, do not use these machine types for new workloads. Ensure that you convert all workloads that use the QEMU **i440fx** machine type to the Q35 machine type before you upgrade to RHOSP 18.0, which requires VM downtime. Bug fixes and support are provided while the feature is deprecated, but Red Hat will not implement new feature enhancements.

BZ#2084206

The use of the QPID Dispatch Router (QDR) for transport of RHOSP telemetry towards Service Telemetry Framework (STF) is deprecated in RHOSP 17.0.

BZ#2090811

The metrics data storage service (gnocchi) has been deprecated since RHOSP 15. Gnocchi is fully supported for storage of metrics when used with the autoscaling use case. For a supported monitoring solution for RHOSP, see [Service Telemetry Framework \(STF\)](#). Use of gnocchi for telemetry storage as a general monitoring solution is not supported.

BZ#2090812

The Alarming service (aodh) has been deprecated since Red Hat OpenStack Platform (RHOSP) 15. The Alarming service is fully supported for delivery of alarms when you use it with the autoscaling use case. For delivery of metrics-based alarms for RHOSP, see [Service Telemetry Framework \(STF\)](#). Use of the Alarming service as part of a general monitoring solution is not supported.

BZ#2100222

The snmp service was introduced to allow the data collection service (Ceilometer) on the undercloud to gather metrics via the snmpd daemon deployed to the overcloud nodes. Telemetry services were previously removed from the undercloud, so the snmp service is no longer necessary or usable in the current state.

BZ#2103869

The Derived Parameters feature is deprecated. It will be removed in a future release. The Derived Parameters feature is configured using the `--plan-environment-file` option of the `openstack overcloud deploy` command.

Workaround / Migration Instructions

HCI overclouds require system tuning. There are many different options for system tuning. The Derived Parameters functionality tuned systems with director by using hardware inspection data and set tuning parameters using the `--plan-environment-file` option of the `openstack overcloud deploy` command. The Derived Parameters functionality is deprecated in Release 17.0 and is removed in 17.1.

The following parameters were tuned by this functionality:

- IsolCpusList
- KernelArgs
- NeutronPhysnetNUMANodesMapping

- NeutronTunnelNUMANodes
- NovaCPUAllocationRatio
- NovaComputeCpuDedicatedSet
- NovaComputeCpuSharedSet
- NovaReservedHostMemory
- OvsDpdkCoreList
- OvsDpdkSocketMemory
- OvsPmdCoreList

To set and tune these parameters starting in 17.0, observe their values using the available command line tools and set them using a standard heat template.

BZ#2128697

The ML2/OVS mechanism driver is deprecated in 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 Red Hat OpenStack Platform (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 will be supported in RHOSP 17.1. Migration tools are available in RHOSP 17.0 for test purposes only.

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](#).

3.1.8. Removed Functionality

BZ#1918403

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

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

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

BZ#1966898

In Red Hat OpenStack Platform 17.0, panko and its API were removed from the distribution.

BZ#1984889

In this release, Block Storage service (cinder) backup support for Google Cloud Services (GCS) has been removed due to a reliance on libraries that are not FIPS compliant.

BZ#2022714

In Red Hat OpenStack Platform 17.0, the collectd-write_redis plugin was removed.

BZ#2023893

In Red Hat OpenStack Platform 17.0, a dependency has been removed from the distribution so that the subpackage **collectd-memcached** cannot be built anymore. The collectd- **memcached** plugin provides similar functionality to that of **collectd-memcached**.

BZ#2065540

In Red Hat OpenStack Platform 17.0, the ability to deliver metrics from collectd to gnocchi was removed.

BZ#2094409

In Red Hat OpenStack Platform 17.0, the deprecated **dbi** and **notify_email** collectd plugins were removed.

BZ#2101948

In Red Hat OpenStack Platform 17.0, the collectd processes plugin has been removed from the default list of plugins. Loading the collectd processes plugin can cause logs to flood with messages, such as "procs_running not found".

BZ#2127184

In Red Hat OpenStack Platform 17.0, support for POWER (ppc64le) architectures has been removed. Only the x86_64 architecture is supported.

3.2. RED HAT OPENSTACK PLATFORM 17.0.1 MAINTENANCE RELEASE - JANUARY 25, 2023

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

3.2.1. Advisory list

This release includes the following advisories:

RHBA-2023:0271

Red Hat OpenStack Platform 17.0.1 bug fix and enhancement advisory

RHBA-2023:0277

Red Hat OpenStack Platform 17.0.1 director images

RHBA-2023:0278

Red Hat OpenStack Platform 17.0.1 director image RPMs

RHBA-2023:0279

Updated Red Hat OpenStack Platform 17.0.1 container images

RHSA-2023:0274

Moderate: Red Hat OpenStack Platform 17.0 (python-XStatic-Angular) security update

RHSA-2023:0275

Moderate: Red Hat OpenStack Platform 17.0 (openstack-neutron) security update

[RHSA-2023:0276](#)

Moderate: Red Hat OpenStack Platform 17.0 (python-ssciclient) security update

3.2.2. Bug Fix

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

BZ#2085583

Before this update, **ovsdb** connection time-outs caused the **nova-compute** agent to become unresponsive. With this update, the issue has been fixed.

BZ#2091076

Before this update, unavailability of the Podman log content caused the health check status script to fail. With this update, an update to the health check status script resolves the issue by using the Podman socket instead of the Podman log. As a result, API health checks, provided through sensibility for Service Telemetry Framework, are now operational.

BZ#2106763

Before this update, an underlying RHEL issue caused a known issue with UEFI boot for instances. With this update, the underlying RHEL issue has now been fixed and the UEFI Secure Boot feature for instances is now available.

BZ#2121098

Before this update, in Red Hat OpenStack Platform (RHOSP) 17.0, Networking service (neutron) requests sometimes failed with a **504 Gateway Time-out** if the request was made when the Networking service reconnected to **ovsdb-server**. These reconnections sometimes happened during failovers or through **ovsdb-server** leader transfers during database compaction.

If neutron debugging was enabled, the Networking service rapidly logged a large number of OVSDDB transaction-returned "TRY_AGAIN" DEBUG messages, until the transaction timed out with an exception.

With this update, the reconnection behavior is fixed to handle this condition, with a single retry of the transaction until a successful reconnection.

BZ#2121634

Before this update, the Red Hat OpenStack Platform (RHOSP) DNS service (designate) was unable to start its central process when TLS-everywhere was enabled. This was caused by an inability to connect to Redis over TLS. With this update in RHOSP 17.0.1, this issue has been resolved.

BZ#2122926

Before this update, adding a member without subnet information when the subnet of the member is different than the subnet of the load balancer Virtual IP (VIP) caused the ovn-octavia provider to wrongly use the VIP subnet for the **subnet_id**, which resulted in no error but no connectivity to the member. With this update, a check that the actual IP of the member belongs to the same CIDR that the VIP belongs to when there is no subnet information resolves the issue. If the two IP addresses do not match, the action is rejected, asking for the **subnet_id**.

BZ#2133029

Before this update, the Alarming service (aodh) used a deprecated gnocchi API to aggregate metrics. This resulted in incorrect metric measures of CPU use in the gnocchi results. With this update, use of dynamic aggregation in gnocchi, which supports the ability to make reaggregations of existing metrics and the ability to make and transform metrics as required, resolves the issue. CPU use in gnocchi is computed correctly.

BZ#2135549

Before this update, deploying RHEL 8.6 images in UEFI mode caused a failure when using the `ironic-python-agent` service because the `ironic-python-agent` service did not understand the RHEL 8.6 UEFI boot loader hint file. With this update, you can now deploy RHEL 8.6 in UEFI mode.

BZ#2138046

Before this update, when you used the whole disk image **overcloud-hardened-uefi-full** to boot overcloud nodes, nodes that used the Legacy BIOS boot mode failed to boot because the **lvmid** of the root volume was different to the **lvmid** referenced in **grub.cfg**. With this update, the **virt-sysprep** task to reset the **lvmid** has been disabled, and nodes with Legacy BIOS boot mode can now be booted with the whole disk image.

BZ#2140881

Before this update, the **network_config** schema in the bare-metal provisioning definition did not allow setting the **num_dpdk_interface_rx_queues** parameter, which caused a schema validation error that blocked the bare-metal node provisioning process. With this update, the schema validation error no longer occurs when the `'num_dpdk_interface_rx_queues'` parameter is used.

3.2.3. Known Issues

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

BZ#2058518

There is currently a known issue when the Object Storage service (`swift`) client blocks a Telemetry service (`ceilometer`) user from fetching object details under the condition of the Telemetry service user having inadequate privileges to poll objects from the Object Storage service. Workaround: Associate the **ResellerAdmin** role with the Telemetry service user by using the command **openstack role add --user ceilometer --project service ResellerAdmin**.

BZ#2104979

A known issue in RHOSP 17.0 prevents the default mechanism for selecting the hypervisor fully qualified domain name (FQDN) from being set properly if the **resource_provider_hypervisors** heat parameter is not set. This causes the single root I/O virtualization (SR-IOV) or Open vSwitch (OVS) agent to fail to start.

Workaround: Specify the hypervisor FQDN explicitly in the heat template. The following is an example of setting this parameter for the SRIOV agent:

```
ExtraConfig: neutron::agents::ml2::sriov::resource_provider_hypervisors: "enp7s0f3:%  
{hiera('fqdn_canonical')},enp5s0f0:%{hiera('fqdn_canonical')}".
```

BZ#2105312

There is currently a known issue where the **ovn/ovsdb_probe_interval** value is not configured in the file **ml2_conf.ini** with the value specified by **OVNOvsdbProbeInterval** because a patch required to configure the neutron server based on **OVNOvsdbProbeInterval** is not included in 17.0.1.

Workaround: Deployments that use **OVNOvsdbProbeInterval** must use **ExtraConfig** hooks in the following manner to configure the neutron server:

```
parameter_defaults:
  OVNOvsdbProbeInterval: <probe interval in milliseconds>
  ControllerExtraConfig:
    neutron::config::plugin_ml2_config:
      ovn/ovsdb_probe_interval:
        value: <probe interval in milliseconds>
```

BZ#2107896

There is currently a known issue that causes tuned kernel configurations to not be applied after initial provisioning.

Workaround: You can use the following custom playbook to ensure that the tuned kernel command line arguments are applied. Save the following playbook as **/usr/share/ansible/tripleo-playbooks/cli-overcloud-node-reset-blscfg.yaml** on the undercloud node:

```
- name: Reset BLSCFG of compute node(s) meant for NFV deployments
  hosts: allovercloud
  any_errors_fatal: true
  gather_facts: true

  pre_tasks:
    - name: Wait for provisioned nodes to boot
      wait_for_connection:
        timeout: 600
        delay: 10

  tasks:
    - name: Reset BLSCFG flag in grub file, if it is enabled
      become: true
      lineinfile:
        path: /etc/default/grub
        line: "GRUB_ENABLE_BLSCFG=false"
        regexp: "^GRUB_ENABLE_BLSCFG=.*"
        insertafter: '^GRUB_DISABLE_RECOVERY.*'
```

Configure the role in the node definition file, **overcloud-baremetal-deploy.yaml**, to run the **cli-overcloud-node-reset-blscfg.yaml** playbook before the playbook that sets the **kernelargs**:

```
- name: ComputeOvsDpdkSriov
  count: 2
  hostname_format: computeovsdpdkSriov-%index%
  defaults:
    networks:
      - network: internal_api
        subnet: internal_api_subnet
      - network: tenant
        subnet: tenant_subnet
      - network: storage
        subnet: storage_subnet
    network_config:
      template: /home/stack/osp17_ref/nic-configs/computeovsdpdkSriov.j2
    config_drive:
      cloud_config:
        ssh_pwauth: true
        disable_root: false
      chpasswd:
        list: |-
          root:12345678
        expire: False
  ansible_playbooks:
    - playbook: /usr/share/ansible/tripleo-playbooks/cli-overcloud-node-reset-blscfg.yaml
    - playbook: /usr/share/ansible/tripleo-playbooks/cli-overcloud-node-kernelargs.yaml
  extra_vars:
    reboot_wait_timeout: 600
```

```

kernel_args: 'default_hugepagesz=1GB hugepagesz=1G hugepages=32 iommu=pt
intel_iommu=on isolcpus=1-11,13-23'
tuned_profile: 'cpu-partitioning'
tuned_isolated_cores: '1-11,13-23'
- playbook: /usr/share/ansible/tripleo-playbooks/cli-overcloud-openvswitch-dpdk.yaml
extra_vars:
memory_channels: '4'
lcore: '0,12'
pmd: '1,13,2,14,3,15'
socket_mem: '4096'
disable_emc: false
enable_tso: false
revalidator: ""
handler: ""
pmd_auto_lb: false
pmd_load_threshold: ""
pmd_improvement_threshold: ""
pmd_rebal_interval: ""
nova_postcopy: true

```

BZ#2125159

There is currently a known issue in RHOSP 17.0 where ML2/OVN deployments fail to automatically create DNS records with the RHOSP DNS service (designate) because the required Networking service (neutron) extension, **dns_domain_ports**, is not present. There is currently no workaround. A fix is planned for a future RHOSP release.

BZ#2127965

There is currently a known issue in RHOSP 17.0 where the Free Range Router (FRR) container does not start after the host on which it resides is rebooted. This issue is caused by a missing file in the BGP configuration. Workaround: Create the file, **/etc/tmpfiles.d/run-frr.conf**, and add the following line:

```
d /run/frr 0750 root root - -
```

After you make this change, **tmpfiles** recreates **/run/frr** after each reboot and the FRR container can start.

CHAPTER 4. TECHNICAL NOTES

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

4.1. RHEA-2022:6543 – RELEASE OF COMPONENTS FOR OSP 17.0

Changes to the ceph component:

- There is currently a known issue where the Swift API does not work and returns a 401 error when multiple Controller nodes are deployed and Ceph is enabled. A workaround is available at <https://access.redhat.com/solutions/6970061>. (BZ#2112988)

Changes to the collectd component:

- In Red Hat OpenStack Platform 17.0, the collectd-write_redis plugin was removed. (BZ#2022714)
- In Red Hat OpenStack Platform 17.0, a dependency has been removed from the distribution so that the subpackage **collectd-memcachec** cannot be built anymore. The collectd- **memcached** plugin provides similar functionality to that of **collectd-memcachec**. (BZ#2023893)
- In Red Hat OpenStack Platform 17.0, the deprecated dbi and notify_email collectd plugins were removed. (BZ#2094409)

Changes to the distribution component:

- In Red Hat OpenStack Platform 17.0, panko and its API were removed from the distribution. (BZ#1966898)
- In Red Hat OpenStack Platform 17.0, the ability to deliver metrics from collectd to gnocchi was removed. (BZ#2065540)

Changes to the openstack-cinder component:

- Before this update, if an operator defined a custom value for the **volume:accept_transfer** policy that referred to the project_id of the user making the volume transfer accept request, the request would fail. This update removes a duplicate policy check that incorrectly compared the project_id of the requestor to the project_id associated with the volume before transfer. The check done at the Block Storage API layer will now function as expected. (BZ#2050773)
- With this enhancement, you can view a volume Encryption Key ID using the cinder client command 'cinder --os-volume-api-version 3.64 volume show <volume_name>'. You must specify microversion 3.64 to view the value. (BZ#1904086)
- Before this update, an issue existed with PowerFlex storage-assisted volume migration when volume migration was performed without conversion of volume type in cases where it should have been converted to thin from thick provisioned. With this update, this issue is fixed. (BZ#1883326)
- In this release, Block Storage service (cinder) backup support for Google Cloud Services (GCS) has been removed due to a reliance on libraries that are not FIPS compliant. (BZ#1984889)

Changes to the openstack-designate component:

- Before this update, a misconfiguration of communication parameters between the DNS service

(designate) worker and deployed BIND instances caused Red Hat OpenStack Platform (RHOSP) 17.0 Beta deployments that have more than one Controller node to fail. With this update, this issue has been resolved, and you can now use the DNS service in a deployment with more than one Controller node. (BZ#1374002)

- In Red Hat OpenStack Platform 17.0, Secure RBAC is available for the DNS service (designate) as a technology preview. (BZ#1901687)

Changes to the openstack-ironic component:

- Before this update, Supermicro servers in UEFI mode would reboot from the network instead of from the local hard disk, causing a failed boot. With this update, Ironic sends the correct raw IPMI commands that request UEFI "boot from hard disk." Booting Supermicro nodes in UEFI mode with IPMI now works as expected. (BZ#1888069)
- This enhancement improves the operating performance of the Bare Metal Provisioning service (ironic) to optimize the performance of large workloads. (BZ#1954274)
- Before this update, network interruptions caused a bare metal node's power state to become **None**, and enter the **maintenance** state. This is due to Ironic's connection cache of Redfish node sessions entering a stale state and not being retried. This state cannot be recovered without restarting the Ironic service. With this update, the underlying REST client has been enhanced to return specific error messages. These error messages are used by Ironic to invalidate cached sessions. (BZ#2064019)

Changes to the openstack-ironic-inspector component:

- Before this update, baremetal node introspection failed with an error and did not retry, when the node had a transient lock on it. With this update, you can perform introspection even when the node has a lock. (BZ#1991657)

Changes to the openstack-manila component:

- With this update, the CephFS drivers in the OpenStack Shared File Systems service (manila) are updated so that you can manage provisioning and storage lifecycle operations by using the Ceph Manager API. When you create new file shares, the shares are created in a new format that is quicker for creating, deleting and operations. This transition does not affect pre-existing file shares. (BZ#1767084)
- With this update, you can restore snapshots with the CephFS Native and CephFS with NFS backends of the Shared File Systems service (manila) by creating a new share from a snapshot. (BZ#1699454)

Changes to the openstack-neutron component:

- You can now migrate the mechanism driver to ML2/OVN from an ML2/OVS deployment that uses the iptables_hybrid firewall driver. The existing instances keep using the hybrid plug mechanism after the migration, but security groups are implemented in OVN and there are no iptables rules present on the compute nodes. (BZ#2075038)
- In an ML2/OVS deployment, Open vSwitch (OVS) does not support offloading OpenFlow rules that have the **skb_priority**, **skb_mark**, or output queue fields set. Those fields are needed to provide quality-of-service (QoS) support for virtio ports. If you set a minimum bandwidth rule for a virtio port, the Neutron Open vSwitch agent marks the traffic of this port with a Packet Mark Field. As a result, this traffic cannot be offloaded, and it affects the traffic in other ports. If you set a bandwidth limit rule, all traffic is marked with the

default 0 queue, which means no traffic can be offloaded.

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

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

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

(BZ#2111015)

Changes to the openstack-nova component:

- Before this update, the help text for the **max_disk_devices_to_attach** parameter did not state that **0** is an invalid value. Also, when the **max_disk_devices_to_attach** parameter was set to **0**, the **nova-compute** service started when it should have failed. With this update, the **max_disk_devices_to_attach** parameter help option text states that a value of **0** is invalid, and if **max_disk_devices_to_attach** is set to **0**, the **nova-compute** service will now log an error and fail to start. (BZ#1801931)

Changes to the openstack-octavia component:

- With this update, the Red Hat OpenStack Platform (RHOSP) 17 Octavia amphora image now includes HAProxy 2.4.x as distributed in Red Hat Enterprise Linux (RHEL) 9. This improves the performance of Octavia load balancers; including load balancers using flavors with more than one vCPU core. (BZ#1813560)
- In Red Hat OpenStack Platform 17.0, secure role-based access control (RBAC) is available for the Load-balancing service (octavia) as a technology preview. (BZ#1901686)

Changes to the openstack-tripleo-common component:

- In RHOSP 17.0 you must use Ceph containers based on RHCSv5.2 GA content. (BZ#2111527)

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

- With this update, **cephadm** and **orchestrator** replace **ceph-ansible**. You can use **director** with **cephadm** to deploy the ceph cluster and additional daemons, and use a new `tripleo-ansible` role to configure and enable the Ceph backend. (BZ#1839169)`
- With this update, Red Hat OpenStack Platform director deployed Ceph includes the RGW daemon, replacing the Object Storage service (swift) for object storage. To keep the Object Storage service, use the **cephadm-rbd-only.yaml** file instead of **cephadm.yaml**. (BZ#1758161)
- With this update, you can now use Red Hat OpenStack Platform director to configure the etcd service to use TLS endpoints when deploying TLS-everywhere. (BZ#1848153)
- In Red Hat OpenStack Platform 17.0, the **iscsi** deployment interface has been deprecated. The default deployment interface is now **direct**. Bug fixes and support are provided while the feature is deprecated but Red Hat will not implement new feature enhancements. In a future release, the interface will be removed. (BZ#1874778)

- This enhancement changes the default machine type for each host architecture to Q35 (**pc-q35-rhel9.0.0**) for new Red Hat OpenStack Platform 17.0 deployments. The Q35 machine type provides several benefits and improvements, including live migration of instances between different RHEL 9.x minor releases, and the native PCIe hotplug that is faster than the ACPI hotplug used by the **i440fx** machine type. (BZ#1946956)
- With this update, the default machine type is RHEL9.0-based Q35 **pc-q35-rhel9.0.0**, with the following enhancements:
 - Live migration across RHEL minor releases.
 - Native PCIe hotplug. This is also ACPI-based like the previous i440fx machine type.
 - Intel input-output memory management unit (IOMMU) emulation helps protect guest memory from untrusted devices that are directly assigned to the guest.
 - Faster SATA emulation.
 - Secure boot. (BZ#1946978)
- In Red Hat OpenStack Platform (RHOSP) 17.0 GA, for NIC-partitioned deployments, you can now pass through virtual functions (VFs) to VMs. To pass through VFs, in a heat environment file, you must specify the VF product ID, vendor ID, and the physical function (PF) PCI addresses:

```
NovaPCIPassthrough:
- product_id: "<VF_product_ID>"
  vendor_id: "<vendor_ID>"
  address: "<PF_PCI_addresses>"
  trusted: "true"
```

The PF PCI address parameter supports string and dict mapping. You can specify wildcard characters and use regular expressions when specifying one or more addresses.

Example

```
NovaPCIPassthrough:
- product_id: "0x7b18"
  vendor_id: "0x8086"
  address: "0000:08:00.*"
  trusted: "true"
```

(BZ#1913862)

- Before this update, the collectd smart plugin required the CAP_SYS_RAWIO capability to work. It was not added by default. With this update, you can add the capability to the collectd container and the smart plugin works. When you use the smart plugin, specify the following parameter in an environment file: CollectdContainerAdditionalCapAdd:
 - "CAP_SYS_RAWIO" (BZ#1984556)
- In Red Hat OpenStack Platform 17.0, the collectd processes plugin has been removed from the default list of plugins. Loading the collectd processes plugin can cause logs to flood with messages, such as "procs_running not found". (BZ#2101948)
- In Red Hat OpenStack Platform (RHOSP) 17.0 GA, a technology preview is available for

integration between the RHOSP Networking service (neutron) ML2/OVN and the RHOSP DNS service (designate). As a result, the DNS service does not automatically add DNS entries for newly created VMs. (BZ#1884782)

Changes to the openstack-tripleo-validations component:

- There is currently a known issue where 'undercloud-heat-purge-deleted' validation fails. This is because it is not compatible with Red Hat OpenStack Platform 17. Workaround: Skip 'undercloud-heat-purge-deleted' with '--skip-list' to skip this validation. (BZ#2105291)

Changes to the puppet-collectd component:

- With this enhancement you can use the PluginInstanceFormat parameter for collectd to specify more than one value. (BZ#1954103)

Changes to the python-octaviaclient component:

- This enhancement includes Octavia support for object tags. This allows users to add metadata to load balancer resources and filter query results based on tags. (BZ#1813573)

Changes to the python-openstackclient component:

- This enhancement includes OpenStack CLI (OSC) support for Block Storage service (cinder) API 3.42. This allows OSC to extend an online volume. (BZ#1689706)

Changes to the python-validations-libs component:

- This enhancement adds the '--limit' argument to the 'openstack tripleo validator show history' command. You can use this argument to show only a specified number of the most recent validations. (BZ#1944872)
- With this update, the Validation Framework provides a configuration file in which you can set parameters for particular use. You can find an example of this file at the root of the code source or in the default location: **/etc/validation.cfg**. You can use the default file in **/etc/** or use your own file and provide it to the CLI with the argument **--config**.

When you use a configuration file there is an order for the variables precedence. The following order is the order of variable precedence:

- User's CLI arguments
- Configuration file
- Default interval values (BZ#1971607)
- With this update, you can supply a new argument **--skiplist** to the **validation run** command. Use this command with a **yaml** file containing services to skip when running validations. (BZ#2013120)

Changes to the tripleo-ansible component:

- This security enhancement reduces the user privilege level required by the OpenStack Shared File System service (manila). You no longer need permissions to create and manipulate Ceph users, because the Shared File Systems service now uses the APIs exposed by the **Ceph Manager** service for this purpose. (BZ#1973356)

- You can now pre-provision bare metal nodes in your application by using the **overcloud node [un]provision** command. (BZ#2041429)
- With this fix, traffic is distributed on VLAN provider networks in ML2/OVN deployments. Previously, traffic on VLAN provider networks was centralized even with the Distributed Virtual Router (DVR) feature enabled. (BZ#2101937)

Changes to the validations-common component:

- This update fixes a bug that incorrectly redirected registered non-stdout callback output from various Ansible processes to the validations logging directory. Output of other processes is no longer stored in validations logging directory. VF callbacks no longer receive information about plays, unless requested. (BZ#1944586)

4.2. RHBA-2023:0271 – RED HAT OPENSTACK PLATFORM 17.0.1 BUG FIX AND ENHANCEMENT ADVISORY

Changes to the openstack-aodh component:

- Before this update, the Alarming service (aodh) used a deprecated gnocchi API to aggregate metrics. This resulted in incorrect metric measures of CPU use in the gnocchi results. With this update, use of dynamic aggregation in gnocchi, which supports the ability to make reaggregations of existing metrics and the ability to make and transform metrics as required, resolves the issue. CPU use in gnocchi is computed correctly. (BZ#2133029)

Changes to the openstack-designate component:

- Before this update, the Red Hat OpenStack Platform (RHOSP) DNS service (designate) was unable to start its central process when TLS-everywhere was enabled. This was caused by an inability to connect to Redis over TLS. With this update in RHOSP 17.0.1, this issue has been resolved. (BZ#2121634)

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

- Before this update, deploying RHEL 8.6 images in UEFI mode caused a failure when using the ironic-python-agent because the ironic-python-agent service did not understand the RHEL 8.6 UEFI boot loader hint file. With this update, you can now deploy RHEL 8.6 in UEFI mode. (BZ#2135549)

Changes to the openstack-nova component:

- Before this update, an underlying RHEL issue caused a known issue with UEFI boot for instances. With this update, the underlying RHEL issue has now been fixed and the UEFI Secure Boot feature for instances is now available. (BZ#2106763)

Changes to the openstack-octavia component:

- Before this update, a race condition occurred in Octavia that may have caused OVN provider load balancers to become stuck in PENDING DELETE under certain conditions. This caused the load balancer to be immutable and unable to update. With this update, the race condition is fixed to resolve the issue. (BZ#2123658)

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

- Before this update, unavailability of the Podman log content caused the health check status script to fail. With this update, an update to the health check status script resolves the issue by

using the Podman socket instead of the Podman log. As a result, API health checks, provided through sensibility for Service Telemetry Framework, are now operational. (BZ#2091076)

- There is currently a known issue in RHOSP 17.0 where the Free Range Router (FRR) container does not start after the host on which it resides is rebooted. This issue is caused by a missing file in the BGP configuration. Workaround: Create the file, `/etc/tmpfiles.d/run-frr.conf`, and add the following line:

```
d /run/frr 0750 root root - -
```

After you make this change, `tmpfiles` recreates `/run/frr` after each reboot and the FRR container can start. (BZ#2127965)

Changes to the python-os-vif component:

- Before this update, `ovsdb` connection time-outs caused the `nova-compute` agent to become unresponsive. With this update, the issue has been fixed. (BZ#2085583)

Changes to the python-ovn-octavia-provider component:

- Before this update, adding a member without subnet information when the subnet of the member is different than the subnet of the load balancer VIP caused the ovn-octavia provider to wrongly use the VIP subnet for the `subnet_id` which resulted in no error but no connectivity to the member. With this update, a check that the actual IP of the member belongs to the same CIDR that the VIP belongs to when there is no subnet information resolves the issue. If the two IP addresses do not match, the action is rejected, asking for the `subnet_id`. (BZ#2122926)
- Before this update, if an ovn-lb is created (VIP and members) in a LS (neutron network) that has 2 subnets (IPv4 and IPv6), and this LS is connected to a LR, removing the LS from the LR leads to the removal of the ovn-lb from the LS and consequently to remove it from the OVN SB DB as it is not associated to any datapath. When re-adding the LS to the LR (the network and subnets to the router) the ovn-lb will not be properly associated to the LR/LS at OVN level and there will be no connectivity
With this update the IP version is checked so that router ports that belong to other subnets are not considered and the ovn-lb is not removed from the LS. This results in the ovn-lb having proper connectivity when a subnet is removed from the router. This resolves the issue. (BZ#2135270)

Changes to the tripleo-ansible component:

- There is currently a known issue that causes tuned kernel configurations to not be applied after initial provisioning.
Workaround: You can use the following custom playbook to ensure that the tuned kernel command line arguments are applied. Save the following playbook as `/usr/share/ansible/tripleo-playbooks/cli-overcloud-node-reset-blscfg.yaml` on the undercloud node:

```
- name: Reset BLSCFG of compute node(s) meant for NFV deployments
  hosts: allovercloud
  any_errors_fatal: true
  gather_facts: true

pre_tasks:
  - name: Wait for provisioned nodes to boot
    wait_for_connection:
      timeout: 600
```

```
delay: 10
```

```
tasks:
```

- name: Reset BLSCFG flag in grub file, if it is enabled
 - become: true
 - lineinfile:
 - path: /etc/default/grub
 - line: "GRUB_ENABLE_BLSCFG=false"
 - regexp: "^GRUB_ENABLE_BLSCFG=."
 - insertafter: '^GRUB_DISABLE_RECOVERY.*'

Configure the role in the node definition file, **overcloud-baremetal-deploy.yaml**, to run the **cli-overcloud-node-reset-blscfg.yaml** playbook before the playbook that sets the **kernelargs**:

```
- name: ComputeOvsDpdkSriov
  count: 2
  hostname_format: computeovsdpdkstriov-%index%
  defaults:
    networks:
      - network: internal_api
        subnet: internal_api_subnet
      - network: tenant
        subnet: tenant_subnet
      - network: storage
        subnet: storage_subnet
    network_config:
      template: /home/stack/osp17_ref/nic-configs/computeovsdpdkstriov.j2
    config_drive:
      cloud_config:
        ssh_pwauth: true
        disable_root: false
      chpasswd:
        list: |-
          root:12345678
        expire: False
    ansible_playbooks:
      - playbook: /usr/share/ansible/tripleo-playbooks/cli-overcloud-node-reset-blscfg.yaml
      - playbook: /usr/share/ansible/tripleo-playbooks/cli-overcloud-node-kernelargs.yaml
    extra_vars:
      reboot_wait_timeout: 600
      kernel_args: 'default_hugepagesz=1GB hugepagesz=1G hugepages=32 iommu=pt
intel_iommu=on isolcpus=1-11,13-23'
      tuned_profile: 'cpu-partitioning'
      tuned_isolated_cores: '1-11,13-23'
      - playbook: /usr/share/ansible/tripleo-playbooks/cli-overcloud-openvswitch-dpdk.yaml
    extra_vars:
      memory_channels: '4'
      lcore: '0,12'
      pmd: '1,13,2,14,3,15'
      socket_mem: '4096'
      disable_emc: false
      enable_tso: false
      revalidator: ""
      handler: ""
      pmd_auto_lb: false
      pmd_load_threshold: ""
```

```
pmd_improvement_threshold: "  
pmd_rebal_interval: "  
nova_postcopy: true
```

(BZ#2107896)

- Before this update, the **network_config** schema in the Bare Metal provisioning definition did not allow setting the **num_dpdk_interface_rx_queues** parameter which caused a schema validation error that blocked the Bare Metal node provisioning process. With this update, the schema validation error no longer occurs when the 'num_dpdk_interface_rx_queues' parameter is used. This resolves the issue. (BZ#2140881)

CHAPTER 5. DOCUMENTATION CHANGES

This section details the major documentation updates delivered with Red Hat OpenStack Platform (RHOSP) 17.0, 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 5.1. Table legend

Column	Meaning
Date	The date that the documentation change was published.
17.0 versions impacted	The RHOSP 17.0 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 5.2. Document changes

Date	17.0 versions impacted	Components	Affected contents	Description of change
20 October 2023	17.0	Networking	Exporting the DNS service pool configuration	Updated the procedure to describe how to run the command inside a container.
04 October 2023	17.0	Networking	<ul style="list-style-type: none"> • Enabling custom composable networks • Removing an overcloud stack 	Replaced networks definition file, network_data.yaml , with network_data_v2.yaml .
29 September 2023	17.0	Security	<ul style="list-style-type: none"> • Federal Information Processing Standard on Red Hat OpenStack Platform 	Corrected procedure so that FIPS images are uploaded to glance
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

Date	17.0 versions impacted	Components	Affected content	Description of change
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.
30 August 2023	17.1	Networking	Overview of allowed address pairs	Added a definition for a virtual port (vport).
30 August 2023	17.0	Security	Creating images	Removed deprecated example for building images and replaced with link to image builder documentation
10 August 2023	17.0	Security	link:https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/17.0/html/users_and_identity_management_guide/assembly_application-credentials#proc_replacing-application-credentials_application-credentials	Procedure on replacing applications credentials in undercloud.conf is rewritten to specify need for user credentials, provides more details.
07 August 2023	17.0	Security	link:https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/17.0/html/security_and_hardening_guide/	Chapter 'Rotating service account passwords' uses a deprecated mistral workflow for execution, and has been removed.
07 August 2023	17.0	Security	Replacing Application Credentials.	Procedure requires use of --unrestricted flag which is not recommended; procedure is removed.
20 July 2023	17.0	All-in-One	Deploying the all-in-one Red Hat OpenStack Platform environment	Procedure is updated with corrected path to clouds.yaml .
12 July 2023	17.0	Security	Implementing TLS-e with Ansible	This procedure is updated with an optional step that is needed when the IdM domain and IDM realm do not match

Date	17.0 versions impacted	Components	Affected content	Description of change
27 June 2023	17.0	Edge	<ul style="list-style-type: none"> link:https://access.redhat.com/documentation/en-us/red_hat_openshift_platform/17.0/html/distributed_compute_node_and_storage_deployment/precaching-glance-images-into-nova#proc_ansible_image_cache_playbook_precaching-glance-images 	The procedure is updated to remove the deprecated way of producing an ansible inventory in Red Hat OpenStack Platform
21 June 2023	17.0	Networking	<ul style="list-style-type: none"> Creating Linux bonds 	The example has changed for the "Linux bond set to 802.3ad LACP mode with one VLAN."
20 June 2023	17.0	Networking	<ul style="list-style-type: none"> Adding a new leaf to a spine-leaf deployment 	Example for flat network mappings (step 7) updated.
13 June 2023	17.0	Networking	<ul style="list-style-type: none"> Load-balancing service (octavia) feature support matrix 	Items added that specify no support for SR-IOV and DPDK.
25 May 2023	17.0	Networking	<ul style="list-style-type: none"> Adding a composable network 	Removed what was previously labeled step 9, assigning predictable virtual IPs for Redis and OVNDBs.
23 May 2023	17.0	Networking	<ul style="list-style-type: none"> Configuring existing BIND servers for the DNS service 	Labeled the feature integrating the RHOSP DNS service with an existing BIND infrastructure as technology preview.
11 May 2023	17.0	Networking	<ul style="list-style-type: none"> Enabling VLAN transparency in ML2/OVN deployments 	Added a new step that instructs users to set --allowed-address on the VM port.

Date	17.0 versions impacted	Components	Affected content	Description of change
4 May 2023	17.0	Validation Framework	<ul style="list-style-type: none">Starting the undercloud in the <i>High Availability for Compute Instances</i> guideStarting the undercloud in the <i>Director Installation and Usage</i> guide	Step 4 has been changed. The option --group pre-introspection has been added.
26 April 2023	17.0	Validation framework	<ul style="list-style-type: none">Running validation using the validation framework	Updated Ansible inventory location and Ansible commands.
19 April 2023	17.0	Networking	<ul style="list-style-type: none">Networking guideIntroduction to the OpenStack DashboardUsing Designate for DNS-as-a-Service	Added "/puppet-generated" to various configuration file paths.

Date	17.0 versions impacted	Components	Affected content	Description of change
18 April 2023	17.0	Compute	<i>Configuring NVDIMM Compute nodes to provide persistent memory for instances</i>	<p>The "Configuring NVDIMM Compute nodes to provide persistent memory for instances" content has been removed from the <i>Configuring the Compute Service for Instance Creation</i> guide. Red Hat has removed support for persistent memory from RHOSP 17.0 and future releases in response to the announcement by the Intel Corporation on July 28, 2022 that they are discontinuing investment in their Intel® Optane™ business:</p> <ul style="list-style-type: none"> ● Intel® Optane™ Business Update: What Does This Mean for Warranty and Support ● Intel® Product Change Notification #119311-00
12 April 2023	17.0	Compute	<ul style="list-style-type: none"> ● Migration constraints 	Updated the configuration to minimize packet loss when live migrating instances in an ML2/OVS deployment.
10 April 2023	17.0	Security	<ul style="list-style-type: none"> ● Removing services from the overcloud firewall 	Removed invalid parameter/value pair "action: accept" from firewall.yaml in example provided in step 2.

Date	17.0 versions impacted	Components	Affected content	Description of change
05 April 2023	17.0	Storage	<ul style="list-style-type: none"> Chapter 6. Configuring the Shared File Systems service (manila) Chapter 7. Performing operations with the Shared File Systems service (manila) 	The Shared File Systems service (manila) content in the <i>Storage Guide</i> has been reorganized into two separate chapters for configuration and operations.
23 Mar 2023	17.0	Networking	<ul style="list-style-type: none"> Configuring VLAN provider networks 	Step 1 under "Verification steps" has been changed. The --external and --share options have been removed.
23 Mar 2023	17.0	Security and Hardening	<ul style="list-style-type: none"> Adding services to the overcloud firewall 	The example of ~/templates/firewall.yaml is updated.
23 Mar 2023	17.0	Networking	<ul style="list-style-type: none"> Fixing OVN controllers that fail to register on edge sites 	A step has been added to the resolution.
20 Mar 2023	17.0	Security	<ul style="list-style-type: none"> Replacing the IdM server for Red Hat OpenStack Platform with a replica 	A new procedure is added to ensure that critical parameters are validated to avoid future deployment failures.
16 Mar 2023	17.0	Storage	<ul style="list-style-type: none"> Deploying Red Hat Ceph Storage and Red Hat OpenStack Platform together with director 	<i>Deploying the Shared File Systems service with CephFS through NFS</i> has been removed from the Customer Portal and the content has been moved to <i>Deploying Red Hat Ceph Storage and Red Hat OpenStack Platform together with director</i> .

Date	17.0 versions impacted	Components	Affected contents	Description of change
07 Mar 2023	17.0	NFV	<ul style="list-style-type: none"> Network Functions Virtualization Planning and Configuration Guide 	Code snippets that contain the devname parameter have been replaced with the address parameter.
07 Mar 2023	17.0	NFV	<ul style="list-style-type: none"> Configuring trust between virtual and physical functions 	The step that instructs users to "Modify permissions to allow users the capability of creating and updating port bindings" (step 3) has been removed.
06 Mar 2023	17.0	Networking	<ul style="list-style-type: none"> Chapter 13. Configuring distributed virtual routing (DVR) 	The topic "Deploying DVR with ML2 OVS" has been removed from the <i>Networking Guide</i> .
02 Mar 2023	17.0	Storage	<ul style="list-style-type: none"> Deploying Red Hat Ceph Storage and Red Hat OpenStack Platform together with director 	<i>Deploying the Shared File Systems service with native CephFS</i> has been removed from the Customer Portal and the content has been moved to <i>Deploying Red Hat Ceph Storage and Red Hat OpenStack Platform together with director</i> .
01 Mar 2023	17.0	Compute and NFV	<ul style="list-style-type: none"> Enabling RT-KVM for NFV Workloads 	Chapter 15, "Configuring real-time compute," has been moved from the <i>Configuring the Compute Service for Instance Creation</i> guide to the <i>Network Functions Virtualization Planning and Configuration Guide</i> .

Date	17.0 versions impacted	Components	Affected content	Description of change
28 Feb 2023	17.0	Hardware Provisioning	<ul style="list-style-type: none"> ● Configuration considerations for overcloud storage nodes ● Provisioning bare metal nodes for the overcloud 	<ul style="list-style-type: none"> ● Added guidance on configuring local disk partition sizes that meet the storage and retention requirements for your storage nodes. ● Added an optional step, step 6, to the "Provisioning bare metal nodes for the overcloud" procedure, to configure the disk partition size allocations if the default disk partition sizes do not meet your requirements.
27 Feb 2023	17.0	Networking	Chapter 4. Configuring the overcloud	Rewrote this chapter to capture deployment changes introduced in RHOSP 17.0.
27 Feb 2023	17.0	Edge	<ul style="list-style-type: none"> ● Updating the central location ● Using external Ceph keys ● Using a pre-installed Red Hat Ceph Storage cluster at the edge 	Fixed the environment file name used within the procedures.
27 Feb 2023	17.0	Hardware Provisioning	Removing failed bare-metal nodes from the node definition file	Added a new procedure on how to remove a failed bare-metal node if the node provisioning fails because of a node hardware or network configuration failure.

Date	17.0 versions impacted	Components	Affected content	Description of change
23 Feb 2023	17.0	Networking, Compute	Enabling VLAN transparency in ML2/OVN deployments	Previously, this procedure instructed you to set the MTU on the network. The updated procedure correctly instructs you to set the MTU on the VLAN interface of each participating VM.
22 Feb 2023	17.0	Networking	Chapter 17. Configuring allowed address pairs	There were several instances of arguments that used underscores (_) instead of hypens (-).
09 Feb 2023	17.0	CloudOps, Storage	Storage Guide	<i>Deployment Recommendations for Specific Red Hat OpenStack Platform Services</i> has been removed from the Customer Portal. For information about recommendations for the Object Storage service (swift), see Configuring the Object Storage service (swift) in the <i>Storage Guide</i> .
08 Feb 2023	17.0	NFV	NFV BIOS settings	Added a note about enabling SR-IOV global and NIC settings in the BIOS.

Date	17.0 versions impacted	Components	Affected content	Description of change
02 Feb 2023	17.0.1	Compute	<ul style="list-style-type: none"> Creating an image for UEFI Secure Boot Image configuration parameters Flavor metadata 	<p>Added content for the UEFI Secure Boot feature:</p> <ul style="list-style-type: none"> New procedure for creating an image for UEFI Secure Boot. Image properties required for a UEFI Secure Boot image: os_secure_boot, hw_firmware_type, and hw_machine_type. Flavor metadata property to enable Secure Boot for instances launched with this flavor: os:secure_boot.
31 Jan 2023	17.0	Networking	Configuring Network service availability zones with ML2/OVN	Included new RHOSP heat parameter, OVNAvailabilityZone .
31 Jan 2023	17.0	NFV	Supported Configurations for NFV Deployments and Chapter 7. Planning your OVS-DPDK deployment	Note added stating that a Support Exception from Red Hat Support is needed to use OVS-DPDK on non-NFV workloads.
27 Jan 2023	17.0	Updates	Chapter 1. Preparing for a minor update	Removed topic about EUS repositories.
25 Jan 2023	17.0	Storage	Volume allocation on multiple back ends	The former Block Storage topic: Specifying back ends for volume creation, has been replaced with: Volume allocation on multiple back ends.
25 Jan 2023	17.0	Updates	Chapter 3. Updating the overcloud	Removed bullet point from list in Section 3.3 due to software fix.

Date	17.0 versions impacted	Components	Affected content	Description of change
23 Jan 2023	17.0	Networking	Chapter 20. Using availability zones to make network resources highly available	Several changes made to identify the distributed compute node (DCN) use case.
17 Jan 2023	17.0	Networking	Chapter 2. Working with ML2/OVN	Two topics have been added to Chapter 2: "Deploying a custom role with ML2/OVN" and "SR-IOV with ML2/OVN and native OVN DHCP."
17 Jan 2023	17.0	Edge	Deploy the edge without storage	Removed redundant step from procedure to deploy storage at the edge
16 Jan 2023	17.0	Networking	Configuring a DHCP relay	Added an important admonition about requiring option 79 for some DHCP relays.
13 Jan 2023	17.0	Edge	Deploying storage at the edge	Replaced instances of deprecated file dcn-hci.yaml with dcn-storage.yaml
13 Jan 2023	17.0	Edge	Deploying storage at the edge	Including the necessary deployed_ceph.yaml and central_ceph_external.yaml in example deploy command.
13 Jan 2023	17.0	Edge	Using a pre-installed Red Hat Ceph Storage cluster at the edge	Changed output directory of openstack overcloud export ceph to for consistency across guide
11 Jan 2023	17.0	Edge	Installing the central location	Fixed ceph deployment to include --stack central parameter and value
22 Dec 2023	17.0	Network Functions Virtualization	Network Functions Virtualization Product Guide	Removed guide because RHOSP 17.0 does not support Network Functions Virtualization (NFV).

Date	17.0 versions impacted	Components	Affected content	Description of change
22 Dec 2023	17.0	Network Functions Virtualization	Network Functions Virtualization Planning and Configuration Guide	Removed guide because RHOSP 17.0 does not support NFV.
22 Dec 2022	17.0	Security	Hardening infrastructure and virtualization	Added procedures for investigating and modifying containers
22 Dec 2022	17.0	Security	Increasing the size of private keys	Added procedure for increasing the default size of private keys
21 Dec 2022	17.0	Compute	Flavor metadata	Added a note about considering the underlying host OS when you use the quota:cpu_* extra specs to tune the instance CPU resource use limits.
20 Dec 2022	17.0	Networking	Configuring Network service availability zones with ML2/OVN	Changes have been made to steps 4 and 5.
20 Dec 2022	17.0	Edge	Deploying edge nodes without storage	Step two is updated, you are not required to generate the DistributedComputeScaleOut role
09 Dec 2022	17.0	Networking	Configuring the Networking service for QoS policies	The step (7.ii.) about resource provider hypervisors has changed.
08 Dec 2022	17.0	Networking	OVN metadata agent on Compute nodes	The corresponding OVN metadata namespace for Virtual Machine (VM) instances on Compute nodes has changed from ovnmeta-<code><datapath_uuid></code> to ovnmeta-<code><network_uuid></code> .
08 Dec 2022	17.0	Networking	QoS rules	A footnote was added to Table 9.1 stating that ML2/OVN does not support DSCP marking QoS policies on tunneled protocols.

Date	17.0 versions impacted	Components	Affected contents	Description of change
30 Nov 2022	17.0	Networking	Configuring Load-balancing service flavors	The three topics in Chapter 6, "Configuring Load-balancing service flavors," erroneously instructed users to access the undercloud to run certain OpenStack commands. Instead, users should access the overcloud.
30 Nov 2022	17.0	Security	Authenticating with keystone	Added procedure to stop repeated failed logins
23 Nov 2022	17.0	Hardware Provisioning	Provisioning bare metal nodes for the overcloud	Updated the guidance on how to configure the href image property in the node definition file.
22 Nov 2022	17.0	Storage	Creating and Managing Images	Updated procedures to use the Image service (glance) command-line client instead of the Dashboard service (horizon) to create and manage images.
9 Nov 2022	17.0	Updates	Performing a minor update of a containerized undercloud	Updated the dnf update command from \$ sudo dnf update -y python3-tripleoclient* ansible to \$ sudo dnf update -y python3-tripleoclient ansible-*
7 Nov 2022	17.0	Updates	Running the overcloud update preparation	Added a prerequisite to regenerate custom NIC templates.
28 Oct 2022	17.0	Backup and Restore	<ul style="list-style-type: none"> ● Installing ReaR on the undercloud node ● Installing ReaR on the control plane nodes 	Updated the command that you use to extract the static ansible inventory file.

Date	17.0 versions impacted	Components	Affected content	Description of change
20 Oct 2022	17.0	Compute	Configuring filters and weights for the Compute scheduler service	Updated the NovaSchedulerDefaultFilters parameter to NovaSchedulerEnabledFilters .
19 Oct 2022	17.0	DCN	Configuring routed spine-leaf in the undercloud	Added procedure for configuring spine/leaf networking on the undercloud.
19 Oct 2022	17.0	DCN	Replacing DistributedComputeHCI nodes	Added procedure for replacing a DCN node.
19 Oct 2022	17.0	Validation	Director Installation and Usage guide	Replaced tripleo validation commands with the new CLI validation commands.
19 Oct 2022	17.0	Validation	Creating a validation	Added procedural content about creating a validation.
19 Oct 2022	17.0	Validation	Changing the validation configuration file	Added procedural content about changing the validation configuration file.
14 Oct 2022	17.0	Identity	Users and Identity Management Guide	Added procedural content about changing the default region name.
14 Oct 2022	17.0	Identity	Users and Identity Management Guide	Added conceptual information about resource credential files.
14 Oct 2022	17.0	Hardware Provisioning	Provisioning and deploying your overcloud	Updated the provisioning step to include details on how to use your own templates instead of the default templates when provisioning the network resources for your physical networks, and when provisioning your bare metal nodes.

Date	17.0 versions impacted	Components	Affected content	Description of change
11 Oct 2022	17.0	Networking	Configuring bridge mappings	Two steps have been added to this procedure that enable customers to change the network name from the default, datacentre .
04 Oct 2022	17.0	Networking	Configuring the Networking service for QoS policies	The example for the SRIOV agent has changed in the in the <i>Networking Guide</i> topic, "Configuring the Networking service for QoS policies."
03 Oct 2022	17.0	Networking	Network definition file configuration options	The default value for mtu has been corrected in the <i>Director Installation and Usage</i> guide topic, "Network definition file configuration options."
30 Sep 2022	17.0	Networking	Cleaning up after Controller node replacement	The note about "bugs prevent the removal of the OVN controller and metadata agents" has been deleted from the <i>Director Installation and Usage</i> guide topic, "Cleaning up after Controller node replacement."
28 Sep 2022	17.0	All	All	In Red Hat OpenStack Platform (RHOSP) 17.0, the heat-admin user has been replaced with the tripleo-admin user.
28 Sep 2022	17.0	Networking	Chapter 6. Troubleshooting networks	Significant changes have been made to the "Troubleshooting networks" chapter in the <i>Networking Guide</i> .
21 Sep 2022	17.0	Networking	Using Designate for DNS-as-a-Service	In Red Hat OpenStack Platform (RHOSP) 17.0, a guide has been added to support the new RHOSP DNS service (designate).

Date	17.0 versions impacted	Components	Affected content	Description of change
21 Sep 2022	17.0	Upgrades	<i>Framework for Upgrades</i> guide	The <i>Framework for Upgrades</i> guide is not published in the RHOSP 17.0 life cycle because upgrades from previous versions are not supported. Upgrades will be supported in RHOSP 17.1 and the <i>Framework for Upgrades Guide</i> will be published. Updates from 17.0.0 to 17.0.z are supported in the RHOSP 17.0 life cycle. For more information, see Keeping Red Hat OpenStack Platform Updated .
21 Sep 2022	17.0	Networking	<i>Testing Migration of the Networking Service to the ML2/OVN Mechanism Driver</i> guide	The <i>Migrating the Networking Service to the ML2/OVN Mechanism Driver</i> guide is published with RHOSP 17.0 for ML2/OVN migration testing purposes only under the title <i>Testing Migration of the Networking Service to the ML2/OVN Mechanism Driver</i> . ML2/OVN migrations are not supported in RHOSP 17.0, because they are not needed for production. Red Hat does not support upgrades to RHOSP 17.0, and all RHOSP 17.0 deployments use the default ML2/OVN mechanism driver. Thus all RHOSP 17.0 deployments start with ML2/OVN and migration is not needed for production.

Date	17.0 versions impacted	Components	Affected content	Description of change
21 Sep 2022	17.0	Compute	<i>Scaling Deployments with Compute Cells</i> guide	The <i>Scaling Deployments with Compute Cells</i> guide is not published for RHOSP 17.0 because the Compute cells feature does not work in RHOSP 17.0. Therefore, the <i>Scaling Deployments with Compute Cells</i> guide has been removed until the underlying issues are fixed.
21 Sep 2022	17.0	All	<ul style="list-style-type: none"> ● Director Installation and Usage guide ● Creating and Managing Images guide ● Storage Guide ● Transitioning to Containerized Services guide ● Security and Hardening Guide 	The <i>Advanced Overcloud Customization</i> guide has been removed for RHOSP 17.0 and the content has been moved to several other guides. For instance, several chapters on networking have been moved to the <i>Director Installation and Usage</i> guide, and the chapter "Configuring the image import method and shared staging area" has been moved to the <i>Creating and Managing Images</i> guide.
21 Sep 2022	17.0	Security	<i>Federate with Identity Service</i> guide	The <i>Federate with Identity Service</i> guide has been removed for RHOSP 17.0. Its contents are consolidated in a Red Hat knowledgebase article that is currently under development.
21 Sep 2022	17.0	Security	Security and Hardening Guide	The <i>Deploy Fernet on the Overcloud</i> guide has been removed. For information about working with Fernet keys, see the Security and Hardening Guide .

Date	17.0 versions impacted	Components	Affected contents	Description of change
21 Sep 2022	17.0	All	Product Documentation for Red Hat OpenStack Platform 17.0	The Product Documentation landing page, also known as splash page, has been reorganized. Sections have been renamed, removed, or replaced and the list of titles represents the latest set of titles.
21 Sep 2022	17.0	All	Deploying Red Hat Ceph Storage and Red Hat OpenStack Platform together with director	The <i>Deploying an overcloud with containerized Red Hat Ceph</i> guide is now called <i>Deploying Red Hat Ceph Storage and Red Hat OpenStack Platform together with director</i> . The content in this document has changed to reflect changes in Red Hat Ceph Storage deployment.
21 Sep 2022	17.0	All	Firewall Rules for Red Hat OpenStack Platform	The <i>Firewall Rules for Red Hat OpenStack Platform</i> guide will not be updated or published in RHOSP 17.0. Red Hat plans to update and publish the guide for RHOSP 17.1.