



Red Hat OpenStack Platform 16.2

Keeping Red Hat OpenStack Platform Updated

Performing minor updates of Red Hat OpenStack Platform

Red Hat OpenStack Platform 16.2 Keeping Red Hat OpenStack Platform Updated

Performing minor updates of Red Hat OpenStack Platform

OpenStack Team
rhos-docs@redhat.com

Legal Notice

Copyright © 2024 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

You can perform a minor update of your Red Hat OpenStack Platform (RHOSP) environment to keep it updated with the latest packages and containers.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	3
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	4
CHAPTER 1. PREPARING FOR A MINOR UPDATE	5
Known issues that might block an update	5
Procedure	7
1.1. UPGRADE PATHS FOR LONG LIFE RELEASES	8
1.2. LOCKING THE ENVIRONMENT TO A RED HAT ENTERPRISE LINUX RELEASE	10
1.3. SWITCHING TO TUS REPOSITORIES	11
1.4. UPDATING RED HAT OPENSTACK PLATFORM AND ANSIBLE REPOSITORIES	13
1.5. SETTING THE CONTAINER-TOOLS MODULE VERSION	15
1.6. UPDATING THE CONTAINER IMAGE PREPARATION FILE	16
1.7. UPDATING THE SSL/TLS CONFIGURATION	16
1.8. DISABLING FENCING IN THE OVERCLOUD	17
CHAPTER 2. UPDATING THE UNDERCLOUD	19
Prerequisites	19
2.1. PERFORMING A MINOR UPDATE OF A CONTAINERIZED UNDERCLOUD	19
2.2. UPDATING THE OVERCLOUD IMAGES	19
CHAPTER 3. UPDATING THE OVERCLOUD	21
Prerequisites	21
Procedure	21
3.1. RUNNING THE OVERCLOUD UPDATE PREPARATION	21
3.2. RUNNING THE CONTAINER IMAGE PREPARATION	22
3.3. OPTIONAL: UPDATING THE OVN-CONTROLLER CONTAINER ON ALL OVERCLOUD SERVERS	23
3.4. UPDATING ALL CONTROLLER NODES	24
3.5. UPDATING ALL COMPUTE NODES	24
3.6. UPDATING ALL HCI COMPUTE NODES	25
3.7. UPDATING ALL DISTRIBUTEDCOMPUTEHCI NODES	26
3.8. UPDATING ALL CEPH STORAGE NODES	27
3.9. PERFORMING ONLINE DATABASE UPDATES	28
3.10. FINALIZING THE UPDATE	28
CHAPTER 4. REBOOTING THE OVERCLOUD	30
4.1. REBOOTING CONTROLLER AND COMPOSABLE NODES	30
4.2. REBOOTING A CEPH STORAGE (OSD) CLUSTER	31
4.3. REBOOTING COMPUTE NODES	32

MAKING OPEN SOURCE MORE INCLUSIVE

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

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

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

Providing documentation feedback in Jira

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

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

CHAPTER 1. PREPARING FOR A MINOR UPDATE

Keep your Red Hat OpenStack Platform (RHOSP) 16.2 environment updated with the latest packages and containers.

You can update the following versions:

Old RHOSP Version	New RHOSP Version
Red Hat OpenStack Platform 16.1.z	Red Hat OpenStack Platform 16.2 latest
Red Hat OpenStack Platform 16.2.z	Red Hat OpenStack Platform 16.2 latest

The RHOSP minor update process workflow

You must complete the following steps to update your RHOSP environment:

1. Prepare your environment for the RHOSP minor update.
2. Update the undercloud to the latest OpenStack 16.2.z version.
3. Update the overcloud to the latest OpenStack 16.2.z version.
4. Upgrade all Red Hat Ceph Storage services.
5. Run the convergence command to refresh your overcloud stack.

If you have a multistack infrastructure, update each overcloud stack completely, one at a time. If you have a distributed compute node (DCN) infrastructure, update the overcloud at the central location completely, and then update the overcloud at each edge site, one at a time.

Considerations before you update your RHOSP environment

To help guide you during the update process, consider the following information:

- Red Hat recommends backing up the undercloud and overcloud control planes. For more information about backing up nodes, see [Backing up and restoring the undercloud and control plane nodes](#).
- Familiarize yourself with the known issues that might block an update.
- Familiarize yourself with the possible update and upgrade paths before you begin your update. For more information, see [Section 1.1, "Upgrade paths for long life releases"](#).
- To identify your current maintenance release, run **\$ cat /etc/rhosp-release**. You can also run this command after updating your environment to validate the update.

Known issues that might block an update

Review the following known issues that might affect a successful minor version update.

[Minor update from 16.1 to 16.2.6 occasionally fails on paunch not being able to start container.](#)

During some updates from 16.1 to 16.2.6, the **collectd** container (sensubility) uses more memory than required, which causes a podman-initiated restart. If a podman-initiated restart occurs during an update, the update fails.

If a podman-initiated restart of the **collectd** container occurs during an update, you must disable the **collectd** container, and then enable the **collectd** container after a successful update. For more information about disabling and enabling the **collectd** container, see the following Red Hat Knowledgebase solution [Updates fail because collectd container \(sensubility\) runs OOM](#).

[BZ#1973660 - \(update\) from 16.1 to 16.2 breaks trying to configure the rabbitmq service.](#)

Overcloud nodes that run Pacemaker version **2.0.3-5.el8_2.4** might fail to update successfully because of a race condition that occurs when shutting down the cluster on a node.

If Pacemaker version **2.0.3-5.el8_2.4** is currently installed on any of the overcloud nodes, you must upgrade Pacemaker before you can update the overcloud nodes. For more information, see the following Red Hat Knowledgebase solution [Update from OSP16.1 to OSP16.2 might fail to update certain HA containers](#).

[BZ#1975240 - update from 16.1 to 16.2, when enabling tsx flag, compute node get restarted during update and ping loss occurs](#)

Starting with Red Hat Enterprise Linux (RHEL) version 8.3, support for the Intel Transactional Synchronization Extensions (TSX) feature is disabled by default. This causes issues with instance live migration between hosts in the following migration scenario:

- Migrating from hosts where the TSX kernel argument is enabled to hosts where the TSX kernel argument is disabled.

Live migration can be unsuccessful in Intel hosts that support the TSX feature. For more information about the CPUs that are affected by this issue, see [Affected Configurations](#).

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

[BZ#1872404 - restarting nodes in parallel while maintaining quorum creates an unexpected node shutdown](#)

For nodes that run RHEL 8.4, and are based on composable roles, you must update the **Database** role first before you can update any other role.

[BZ#2027787 - Undercloud upgrade to 16.2 fails because of missing dependencies of swtpm](#)

There is a known issue with the **advanced-virt-for-rhel-8-x86_64-eus-rpms** and **advanced-virt-for-rhel-8-x86_64-rpms** repositories that prevents a successful upgrade. To disable these repositories before upgrading, see the Red Hat Knowledgebase solution [advanced-virt-for-rhel-8-x86_64-rpms are no longer required in OSP 16.2](#).

[BZ#2009106 - podman panic after tripleo_nova_libvirt restart two times](#)

There is a known issue with upgrading from RHOSP 16.1 to 16.2, and with upgrading from RHOSP 16.2.1 to 16.2.2, related to changes in Podman and the libvirt service. If you do not migrate workloads before upgrading, then the upgrade might fail.

[BZ#2129445 - 16.2.0 known issue nova_libvirt tag: 16.2.0-55.1638436404 libvirt version incompatibility leaves instances in unmanageable state after minor update from RHOSP 16.2.0 to RHOSP 16.2.x](#)

Do not update from RHOSP 16.2.0 to 16.2.2 or 16.2.3 until you evaluate your risk of serious impact from a libvirt version incompatibility.

To evaluate your risk, complete the following steps:

1. Check the libvirt package in the **nova_libvirt** container on all Compute nodes:

```
$ sudo podman exec nova_libvirt rpm -qa libvirt-*
```

2. Check the libvirt version of the **nova_compute** container:

```
$ sudo podman exec nova_compute rpm -qa libvirt-*
```

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

If the libvirt version is 7.6, the deployment is affected by the bug. Your update is at risk. To update your deployment, follow the steps in [Workaround for a libvirt version-compat issue \(bug 2109350\) when updating RHOSP 16.2.0](#).

[BZ#2181566 - Overcloud servers on 16.2.3 show short fqdn hostname](#)

In Red Hat OpenStack Platform (RHOSP) 16.2, the **nova::dhcp_domain** parameter was introduced. If you update from RHOSP 16.1 to any 16.2 release, and your custom template includes the legacy **nova::metadata::dhcp_domain** parameter, a conflict occurs with the **nova::dhcp_domain** parameter. As a result, hostnames do not generate on Compute nodes. To avoid this issue, choose one of the following options:

- Set the legacy **nova::metadata::dhcp_domain** and **nova::dhcp_domain** parameters to the same value.
- Wait to update. A fix is planned in RHOSP 16.2.6.

[BZ#2078793 - Unable to delete an LB member after OSP16.1→16.2 upgrade](#)

Do not alter OVN DB entries during an update from RHOSP 16.1 to 16.2, if the update includes an OVN DB schema upgrade. Doing so can result in misconfiguration and data loss.

If you alter the OVN DB during an update that includes an OVN DB schema upgrade and OpenShift, Kuryr, and the Load-balancing service (octavia), you might not be able to delete Load-balancing entities.

Workaround: If you altered the OVN DB during an update that includes an OVN DB schema upgrade and OpenShift, Kuryr, and the Load-balancing service, and you cannot delete Load-balancing entities, perform the following steps:

1. Access the mysql octavia DB.
2. Change the entity's **provisioning_status** to **DELETED**.

If the issue occurs over any other OVN DB entity after altering the OVN DB during an update, run the **neutron-db-sync tool**.

Procedure


To prepare your RHOSP environment for the minor update, complete the following procedures:

1. [Section 1.2, "Locking the environment to a Red Hat Enterprise Linux release"](#)

2. [Section 1.3, "Switching to TUS repositories"](#)
3. [Section 1.4, "Updating Red Hat Openstack Platform and Ansible repositories"](#)
4. [Section 1.5, "Setting the container-tools module version"](#)
5. [Section 1.6, "Updating the container image preparation file"](#)
6. [Section 1.7, "Updating the SSL/TLS configuration"](#)
7. [Section 1.8, "Disabling fencing in the overcloud"](#)

1.1. UPGRADE PATHS FOR LONG LIFE RELEASES

Familiarize yourself with the possible update and upgrade paths before you begin an update or an upgrade.



NOTE

You can view your current RHOSP and RHEL versions in the `/etc/rhosp-release` and `/etc/redhat-release` files.

Table 1.1. Updates version path

Current version	Target version
RHOSP 10.0.x on RHEL 7.x	RHOSP 10.0 latest on RHEL 7.7 latest
RHOSP 13.0.x on RHEL 7.x	RHOSP 13.0 latest on RHEL 7.9 latest
RHOSP 16.1.x on RHEL 8.2	RHOSP 16.1 latest on RHEL 8.2 latest
RHOSP 16.1.x on RHEL 8.2	RHOSP 16.2 latest on RHEL 8.4 latest
RHOSP 16.2.x on RHEL 8.4	RHOSP 16.2 latest on RHEL 8.4 latest

Table 1.2. Upgrades version path

Current version	Target version
RHOSP 10 on RHEL 7.7	RHOSP 13 latest on RHEL 7.9 latest
RHOSP 13 on RHEL 7.9	RHOSP 16.1 latest on RHEL 8.2 latest
RHOSP 13 on RHEL 7.9	RHOSP 16.2 latest on RHEL 8.4 latest

For more information, see [Framework for Upgrades \(13 to 16.2\)](#).

Red Hat provides two options for upgrading your environment to the next long life release:

In-place upgrade

Perform an upgrade of the services in your existing environment. This guide primarily focuses on this option.

Parallel migration

Create a new Red Hat OpenStack Platform 16.2 environment and migrate your workloads from your current environment to the new environment. For more information about Red Hat OpenStack Platform parallel migration, contact Red Hat Global Professional Services.



IMPORTANT

The durations in this table are minimal estimates based on internal testing and might not apply to all production environments. For example, if your hardware has low specifications or an extended boot period, allow for more time with these durations. To accurately gauge the upgrade duration for each task, perform these procedures in a test environment with hardware similar to your production environment.

Table 1.3. Impact and duration of upgrade paths

	In-place upgrade	Parallel migration
Upgrade duration for undercloud	Estimated duration for each major action includes the following: <ul style="list-style-type: none"> ● 30 minutes for Leapp upgrade command ● 30 minutes for Leapp reboot ● 40 minutes for director upgrade 	None. You are creating a new undercloud in addition to your existing undercloud.
Upgrade duration for overcloud control plane	Estimates for each Controller node: <ul style="list-style-type: none"> ● 60 minutes for Leapp upgrade and reboot ● 60 minutes for service upgrade 	None. You are creating a new control plane in addition to your existing control plane.
Outage duration for control plane	The duration of the service upgrade of the bootstrap Controller node, which is approximately 60 minutes.	None. Both overclouds are operational during the workload migration.
Consequences of control plane outage	You cannot perform OpenStack operations during the outage.	No outage.

	In-place upgrade	Parallel migration
Upgrade duration for overcloud data plane	Estimates for each Compute node and Ceph Storage node: <ul style="list-style-type: none"> • 60 minutes for Leapp upgrade and reboot • 30 minutes for service upgrade 	None. You are creating a new data plane in addition to your existing data plane.
Outage duration for data plane	The outage is minimal due to workload migration from node to node.	The outage is minimal due to workload migration from overcloud to overcloud.
Additional hardware requirements	No additional hardware is required.	Additional hardware is required to create a new undercloud and overcloud.

1.2. LOCKING THE ENVIRONMENT TO A RED HAT ENTERPRISE LINUX RELEASE

Red Hat OpenStack Platform (RHOSP) 16.2 is supported on Red Hat Enterprise Linux (RHEL) 8.4. Before you perform the update, lock the undercloud and overcloud repositories to the RHEL 8.4 release to avoid upgrading the operating system to a newer minor release.

Procedure

1. Log in to the undercloud as the **stack** user.
2. Source the **stackrc** file:

```
$ source ~/stackrc
```

3. Edit your overcloud subscription management environment file, which is the file that contains the **RhsmVars** parameter. The default name for this file is usually **rhsm.yml**.
4. Check if your subscription management configuration includes the **rhsm_release** parameter. If the **rhsm_release** parameter is not present, add it and set it to 8.4:

```
parameter_defaults:
  RhsmVars:
    ...
    rhsm_username: "myusername"
    rhsm_password: "p@55w0rd!"
    rhsm_org_id: "1234567"
    rhsm_pool_ids: "1a85f9223e3d5e43013e3d6e8ff506fd"
    rhsm_method: "portal"
    rhsm_release: "8.4"
```

5. Save the overcloud subscription management environment file.
6. Create a static inventory file of your overcloud:

```
$ tripleo-ansible-inventory --ansible_ssh_user heat-admin --static-yaml-inventory
~/inventory.yaml
```

If you use an overcloud name that is different to the default overcloud name of **overcloud**, set the name of your overcloud with the **--plan** option.

7. Create a playbook that contains a task to lock the operating system version to RHEL 8.4 on all nodes:

```
$ cat > ~/set_release.yaml <<'EOF'
- hosts: all
  gather_facts: false
  tasks:
    - name: set release to 8.4
      command: subscription-manager release --set=8.4
      become: true
EOF
```

8. Run the **set_release.yaml** playbook:

```
$ ansible-playbook -i ~/inventory.yaml -f 25 ~/set_release.yaml --limit <undercloud>,
<Controller>,<Compute>
```

- Use the **--limit** option to apply the content to all RHOSP nodes. Replace **<undercloud>**, **<Controller>**, **<Compute>** with the Ansible groups in your environment that contain those nodes.
- You cannot run this playbook against Ceph Storage nodes if you are using a different subscription for these nodes.



NOTE

To manually lock a node to a version, log in to the node and run the **subscription-manager release** command:

```
$ sudo subscription-manager release --set=8.4
```

1.3. SWITCHING TO TUS REPOSITORIES

Your Red Hat OpenStack Platform (RHOSP) subscription includes repositories for Red Hat Enterprise Linux (RHEL) 8.4 Extended Update Support (EUS), in addition to standard repositories. After May 30, 2023, you must enable the RHEL 8.4 Telecommunications Update Service (TUS) repositories for Maintenance Support. The TUS repositories include the latest security patches and bug fixes for RHEL 8.4.

Switch your repositories to the required TUS repositories before you perform an update.

Table 1.4. Switching from EUS repositories to TUS repositories

EUS repository	TUS repository
rhel-8-for-x86_64-baseos-eus-rpms	rhel-8-for-x86_64-baseos-tus-rpms
rhel-8-for-x86_64-appstream-eus-rpms	rhel-8-for-x86_64-appstream-tus-rpms
rhel-8-for-x86_64-highavailability-eus-rpms	rhel-8-for-x86_64-highavailability-tus-rpms

Table 1.5. Switching from standard repositories to TUS repositories

Standard repository	TUS repository
rhel-8-for-x86_64-baseos-rpms	rhel-8-for-x86_64-baseos-tus-rpms
rhel-8-for-x86_64-appstream-rpms	rhel-8-for-x86_64-appstream-tus-rpms
rhel-8-for-x86_64-highavailability-rpms	rhel-8-for-x86_64-highavailability-tus-rpms

**IMPORTANT**

You must use TUS repositories to retain compatibility with a specific version of Podman. Later versions of Podman are untested with Red Hat OpenStack Platform 16.2 and can cause unexpected results.

Procedure

1. Log in to the undercloud as the **stack** user.
2. Source the **stackrc** file:

```
$ source ~/stackrc
```

3. Edit your overcloud subscription management environment file, which is the file that contains the **RhsmVars** parameter. The default name for this file is usually **rhsm.yml**.
4. Check the **rhsm_repos** parameter in your subscription management configuration. If this parameter does not include the TUS repositories, change the relevant repositories to the TUS versions:

```
parameter_defaults:
  RhsmVars:
    rhsm_repos:
      - rhel-8-for-x86_64-baseos-tus-rpms
      - rhel-8-for-x86_64-appstream-tus-rpms
      - rhel-8-for-x86_64-highavailability-tus-rpms
      - ansible-2.9-for-rhel-8-x86_64-rpms
      - openstack-16.2-for-rhel-8-x86_64-rpms
      - rhceph-4-tools-for-rhel-8-x86_64-rpms
      - fast-datapath-for-rhel-8-x86_64-rpms
```


5. Save the overcloud subscription management environment file.

6. Create a static inventory file of your overcloud:

```
$ tripleo-ansible-inventory --ansible_ssh_user heat-admin --static-yaml-inventory
~/inventory.yaml
```

If you use an overcloud name that is different to the default overcloud name of **overcloud**, set the name of your overcloud with the **--plan** option.

7. Create a playbook that contains a task to set the repositories to RHEL 8.4 TUS on all nodes:

```
$ cat > ~/change_tus.yaml <<'EOF'
- hosts: all
  gather_facts: false
  tasks:
    - name: change to tus repos
      command: subscription-manager repos --disable=rhel-8-for-x86_64-baseos-eus-rpms --
disable=rhel-8-for-x86_64-appstream-eus-rpms --disable=rhel-8-for-x86_64-highavailability-
eus-rpms --enable=rhel-8-for-x86_64-baseos-tus-rpms --enable=rhel-8-for-x86_64-
appstream-tus-rpms --enable=rhel-8-for-x86_64-highavailability-tus-rpms
      become: true
EOF
```

- If your environment includes standard repositories, disable the following repositories:
 - rhel-8-for-x86_64-baseos-rpms
 - rhel-8-for-x86_64-appstream-rpms
 - rhel-8-for-x86_64-highavailability-rpms

8. Run the **change_tus.yaml** playbook:

```
$ ansible-playbook -i ~/inventory.yaml -f 25 ~/change_tus.yaml --limit <undercloud>,
<Controller>,<Compute>
```

- Use the **--limit** option to apply the content to all Red Hat OpenStack Platform nodes. Replace **<undercloud>**, **<Controller>**, **<Compute>** with the Ansible groups in your environment that contain those nodes.
- You cannot run this playbook against Ceph Storage nodes if you are using a different subscription for these nodes.

1.4. UPDATING RED HAT OPENSTACK PLATFORM AND ANSIBLE REPOSITORIES

Update your repositories to use Red Hat OpenStack Platform (RHOSP) 16.2 and Ansible 2.9 packages.

Procedure

1. Log in to the undercloud as the **stack** user.
2. Source the **stackrc** file:

```
$ source ~/stackrc
```

3. Edit your overcloud subscription management environment file, which is the file that contains the **RhsmVars** parameter. The default name for this file is usually **rhsm.yml**.
4. Check the **rhsm_repos** parameter in your subscription management configuration. If the **rhsm_repos** parameter uses the RHOSP 16.1 and Ansible 2.8 repositories, change the repository to the correct versions:

```
parameter_defaults:
  RhsmVars:
    rhsm_repos:
      - rhel-8-for-x86_64-baseos-tus-rpms
      - rhel-8-for-x86_64-appstream-tus-rpms
      - rhel-8-for-x86_64-highavailability-tus-rpms
      - ansible-2.9-for-rhel-8-x86_64-rpms
      - openstack-16.2-for-rhel-8-x86_64-rpms
      - fast-datapath-for-rhel-8-x86_64-rpms
```

5. Save the overcloud subscription management environment file.
6. Create a static inventory file of your overcloud:

```
$ tripleo-ansible-inventory --ansible_ssh_user heat-admin --static-yaml-inventory
~/inventory.yaml
```

If you use an overcloud name that is different to the default overcloud name of **overcloud**, set the name of your overcloud with the **--plan** option.

7. Create a playbook that contains a task to set the repositories to RHOSP 16.2 on all RHOSP nodes:

```
$ cat > ~/update_rhosp_repos.yaml <<'EOF'
- hosts: all
  gather_facts: false
  tasks:
    - name: change osp repos
      command: subscription-manager repos --disable=openstack-16.1-for-rhel-8-x86_64-rpms
--enable=openstack-16.2-for-rhel-8-x86_64-rpms --disable=ansible-2.8-for-rhel-8-x86_64-
rpms --enable=ansible-2.9-for-rhel-8-x86_64-rpms
      become: true
EOF
```

8. Run the **update_rhosp_repos.yaml** playbook:

```
$ ansible-playbook -i ~/inventory.yaml -f 25 ~/update_rhosp_repos.yaml --limit <undercloud>,
<Controller>,<Compute>
```

- Use the **--limit** option to apply the content to all RHOSP nodes. Replace **<undercloud>**, **<Controller>**, **<Compute>** with the Ansible groups in your environment that contain those nodes.
- You cannot run this playbook against Ceph Storage nodes if you are using a different subscription for these nodes.

9. Create a playbook that contains a task to set the repositories to RHOSP 16.2 on all Ceph Storage nodes:

```
$ cat > ~/update_ceph_repos.yaml <<'EOF'
- hosts: all
  gather_facts: false
  tasks:
    - name: change ceph repos
      command: subscription-manager repos --disable=openstack-16-deployment-tools-for-
rhel-8-x86_64-rpms --enable=openstack-16.2-deployment-tools-for-rhel-8-x86_64-rpms --
disable=ansible-2.8-for-rhel-8-x86_64-rpms --enable=ansible-2.9-for-rhel-8-x86_64-rpms
      become: true
EOF
```

10. Run the **update_ceph_repos.yaml** playbook:

```
$ ansible-playbook -i ~/inventory.yaml -f 25 ~/update_ceph_repos.yaml --limit CephStorage
```

Use the **--limit** option to apply the content to Ceph Storage nodes.

1.5. SETTING THE CONTAINER-TOOLS MODULE VERSION

Set the **container-tools** module to version **3.0** to ensure that you use the correct package versions on all nodes.

Procedure

1. Log in to the undercloud as the **stack** user.
2. Source the **stackrc** file:

```
$ source ~/stackrc
```

3. Create a static inventory file of your overcloud:

```
$ tripleo-ansible-inventory --ansible_ssh_user heat-admin --static-yaml-inventory
~/inventory.yaml
```

If you use an overcloud name that is different to the default overcloud name of **overcloud**, set the name of your overcloud with the **--plan** option.

4. Create a playbook that contains a task to set the **container-tools** module to version **3.0** on all nodes:

```
$ cat > ~/container-tools.yaml <<'EOF'
- hosts: all
  gather_facts: false
  tasks:
    - name: disable default dnf module for container-tools
      command: dnf module reset -y container-tools
      become: true
    - name: set dnf module for container-tools:3.0
```

```
command: dnf module enable -y container-tools:3.0
become: true
EOF
```

5. Run the **container-tools.yaml** playbook against all nodes:

```
$ ansible-playbook -i ~/inventory.yaml -f 25 ~/container-tools.yaml
```

1.6. UPDATING THE CONTAINER IMAGE PREPARATION FILE


The container preparation file is the file that contains the **ContainerImagePrepare** parameter. You use this file to define the rules for obtaining container images for the undercloud and overcloud.

Before you update your environment, check the file to ensure that you obtain the correct image versions.

Procedure

1. Edit the container preparation file. The default name for this file is usually **containers-prepare-parameter.yaml**.
2. Check the **tag** parameter is set to **16.2** for each rule set:

```
parameter_defaults:
  ContainerImagePrepare:
    - push_destination: true
      set:
        ...
        tag: '16.2'
        tag_from_label: '{version}-{release}'
```



NOTE

If you do not want to use a specific tag for the update, such as **16.2** or **16.2.2**, remove the **tag** key-value pair and specify **tag_from_label** only. This uses the installed Red Hat OpenStack Platform version to determine the value for the tag to use as part of the update process.

3. Save this file.

1.7. UPDATING THE SSL/TLS CONFIGURATION

Remove the **NodeTLSData** resource from the **resource_registry** to update your SSL/TLS configuration.

Procedure

1. Log in to the undercloud as the **stack** user.
2. Source the **stackrc** file:

```
$ source ~/stackrc
```

3. Edit your custom overcloud SSL/TLS public endpoint file, which is usually named `~/templates/enable-tls.yaml`.
4. Remove the **NodeTLSData** resource from the **resource_registry**:

```
resource_registry:
  OS::TripleO::NodeTLSData: /usr/share/openstack-tripleo-heat-
    templates/puppet/extraconfig/tls/tls-cert-inject.yaml
  ...
```

The overcloud deployment uses a new service in HAProxy to determine if SSL/TLS is enabled.



NOTE

If this is the only resource in the **resource_registry** section of the **enable-tls.yaml** file, remove the complete **resource_registry** section.

5. Save the SSL/TLS public endpoint file.
6. If you are updating from Red Hat OpenStack Platform 16.1, you must update the permissions in Red Hat Identity Manager (IdM) for all pre-update checks to pass. Use **ssh** to login to the server that is running your IdM, and then run the following commands:

```
$ kinit admin
$ ipa privilege-add-permission 'Nova Host Management' --permission 'System: Modify Realm Domains'
```

1.8. DISABLING FENCING IN THE OVERCLOUD

Before you update the overcloud, ensure that fencing is disabled.

If fencing is deployed in your environment during the Controller nodes update process, the overcloud might detect certain nodes as disabled and attempt fencing operations, which can cause unintended results.

If you have enabled fencing in the overcloud, you must temporarily disable fencing for the duration of the update to avoid any unintended results.

Procedure

1. Log in to the undercloud as the **stack** user.
2. Source the **stackrc** file.
3. Log in to a Controller node and run the Pacemaker command to disable fencing:

```
$ ssh heat-admin@<controller_ip> "sudo pcs property set stonith-enabled=false"
```

Replace **<controller_ip>** with the IP address of a Controller node. You can find the IP addresses of your Controller nodes with the **openstack server list** command.

4. In the **fencing.yaml** environment file, set the **EnableFencing** parameter to **false** to ensure that fencing stays disabled during the update process.

Additional Resources

- [Fencing Controller nodes with STONITH](#)

CHAPTER 2. UPDATING THE UNDERCLOUD

You can use director to update the main packages on the undercloud node. To update the undercloud and its overcloud images to the latest Red Hat OpenStack Platform (RHOSP) 16.2 version, complete the following procedures:

1. [Section 2.1, “Performing a minor update of a containerized undercloud”](#)
2. [Section 2.2, “Updating the overcloud images”](#)

Prerequisites

- Before you can update the undercloud to the latest RHOSP 16.2 version, ensure that you complete all the update preparation procedures. For more information, see [Chapter 1, *Preparing for a minor update*](#).

2.1. PERFORMING A MINOR UPDATE OF A CONTAINERIZED UNDERCLOUD

Director provides commands to update the main packages on the undercloud node. Use director to perform a minor update within the current version of your RHOSP environment.

Procedure

1. On the undercloud node, log in as the **stack** user.
2. Source the **stackrc** file:

```
$ source ~/stackrc
```

3. Update the director main packages with the **dnf update** command:

```
$ sudo dnf update -y python3-tripleoclient* tripleo-ansible ansible
```

4. Update the undercloud environment with the **openstack undercloud upgrade** command :

```
$ openstack undercloud upgrade
```

5. Wait until the undercloud update process completes.
6. Reboot the undercloud to update the operating system’s kernel and other system packages:

```
$ sudo reboot
```

7. Wait until the node boots.

2.2. UPDATING THE OVERCLOUD IMAGES

You must replace your current overcloud images with new versions to ensure that director can introspect and provision your nodes with the latest version of the RHOSP software. If you are using pre-provisioned nodes, this step is not required.

Prerequisites

- You have updated the undercloud node to the latest version. For more information, see [Section 2.1, “Performing a minor update of a containerized undercloud”](#).

Procedure

- Source the **stackrc** file:

```
$ source ~/stackrc
```

- Remove any existing images from the **images** directory on the **stack** user’s home (**/home/stack/images**):

```
$ rm -rf ~/images/*
```

- Extract the archives:

```
$ cd ~/images
$ for i in /usr/share/rhosp-director-images/overcloud-full-latest-16.2.tar /usr/share/rhosp-director-images/ironic-python-agent-latest-16.2.tar; do tar -xvf $i; done
$ cd ~
```

- Import the latest images into the director:

```
$ openstack overcloud image upload --update-existing --image-path /home/stack/images/
```

- Configure your nodes to use the new images:

```
$ openstack overcloud node configure $(openstack baremetal node list -c UUID -f value)
```

- Verify the existence of the new images:

```
$ openstack image list
$ ls -l /var/lib/ironic/httpboot
```

IMPORTANT

- When you deploy overcloud nodes, ensure that the overcloud image version corresponds to the respective heat template version. For example, use only the RHOSP 16.2 images with the RHOSP 16.2 heat templates.
- If you deployed a connected environment that uses the Red Hat Customer Portal or Red Hat Satellite Server, the overcloud image and package repository versions might be out of sync. To ensure that the overcloud image and package repository versions match, you can use the **virt-customize** tool. For more information, see the Red Hat Knowledgebase solution [Modifying the Red Hat Linux OpenStack Platform Overcloud Image with virt-customize](#).
- The new **overcloud-full** image replaces the old **overcloud-full** image. If you made changes to the old image, you must repeat the changes in the new image, especially if you want to deploy new nodes in the future.

CHAPTER 3. UPDATING THE OVERCLOUD

After you update the undercloud, you can update the overcloud by running the overcloud and container image preparation commands, updating your nodes, and running the **overcloud update converge** command. The control plane API is fully available during a minor update.

Prerequisites

- You have updated the undercloud node to the latest version. For more information, see [Chapter 2, Updating the undercloud](#).
- If you use a local set of core templates in your **stack** user home directory, ensure that you update the templates and use the recommended workflow in [Using Customized Core Heat Templates](#) in the *Advanced Overcloud Customization* guide. You must update the local copy before you upgrade the overcloud.

Procedure

To update the overcloud, you must complete the following procedures:

1. [Section 3.1, "Running the overcloud update preparation"](#)
2. [Section 3.2, "Running the container image preparation"](#)
3. [Section 3.3, "Optional: Updating the ovn-controller container on all overcloud servers"](#)
4. [Section 3.4, "Updating all Controller nodes"](#)
5. [Section 3.5, "Updating all Compute nodes"](#)
6. [Section 3.6, "Updating all HCI Compute nodes"](#)
7. [Section 3.8, "Updating all Ceph Storage nodes"](#)
8. [Section 3.9, "Performing online database updates"](#)
9. [Section 3.10, "Finalizing the update"](#)

3.1. RUNNING THE OVERCLOUD UPDATE PREPARATION

To prepare the overcloud for the update process, you must run the **openstack overcloud update prepare** command, which updates the overcloud plan to Red Hat OpenStack Platform (RHOSP) 16.2 and prepares the nodes for the update.

Prerequisites

- If you use a Ceph subscription and have configured director to use the **overcloud-minimal** image for Ceph storage nodes, you must ensure that in the **roles_data.yaml** role definition file, the **rhsm_enforce** parameter is set to **False**.
- If you rendered custom NIC templates, you must regenerate the templates with the updated version of the **openstack-tripleo-heat-templates** collection to avoid incompatibility with the overcloud version. For more information about custom NIC templates, see [Rendering default network interface templates for customization](#) in the *Advanced Overcloud Customization* guide.

**NOTE**

For distributed compute node (edge) architectures with OVN deployments, you must complete this procedure for each stack with Compute, DistributedCompute, or DistributedComputeHCI nodes before proceeding with section [Updating the ovn-controller container on all overcloud servers](#).

Procedure

1. Source the **stackrc** file:

```
$ source ~/stackrc
```

2. Run the update preparation command:

```
$ openstack overcloud update prepare \
  --templates \
  --stack <stack_name> \
  -r <roles_data_file> \
  -n <network_data_file> \
  -e <environment_file> \
  -e <environment_file> \
  ...
```

Include the following options relevant to your environment:

- If the name of your overcloud stack is different to the default name **overcloud**, include the **--stack** option in the update preparation command and replace **<stack_name>** with the name of your stack.
 - If you use your own custom roles, include your custom roles (**<roles_data>**) file (**-r**).
 - If you use custom networks, include your composable network (**<network_data>**) file (**-n**).
 - If you deploy a high availability cluster, include the **--ntp-server** option in the update preparation command, or include the **NtpServer** parameter and value in your environment file.
 - Any custom configuration environment files (**-e**).
3. Wait until the update preparation process completes.

3.2. RUNNING THE CONTAINER IMAGE PREPARATION

Before you can update the overcloud, you must prepare all container image configurations that are required for your environment and pull the latest RHOSP 16.2 container images to your undercloud.

To complete the container image preparation, you must run the **openstack overcloud external-update run** command against tasks that have the **container_image_prepare** tag.

**NOTE**

If you are not using the default stack name, which is **overcloud**, set your stack name with the **--stack <stack_name>** option and replace **<stack_name>** with the name of your stack.

Procedure

1. Source the **stackrc** file:

```
$ source ~/stackrc
```

2. Run the **openstack overcloud external-update run** command against tasks that have the **container_image_prepare** tag:

```
$ openstack overcloud external-update run --stack <stack_name> --tags
container_image_prepare
```

3.3. OPTIONAL: UPDATING THE OVN-CONTROLLER CONTAINER ON ALL OVERCLOUD SERVERS

If you deployed your overcloud with the Modular Layer 2 Open Virtual Network mechanism driver (ML2/OVN), update the ovn-controller container to the latest RHOSP 16.2 version. The update occurs on every overcloud server that runs the ovn-controller container.



IMPORTANT

The following procedure updates the ovn-controller containers on servers that are assigned the Compute role before it updates the ovn-northd service on servers that are assigned the Controller role.

If you accidentally updated the ovn-northd service before following this procedure, you might not be able to reach your virtual machines or create new virtual machines or virtual networks. The following procedure restores connectivity.



NOTE

For distributed compute node (edge) architectures, you must complete this procedure for each stack with Compute, DistributedCompute, or DistributedComputeHCI nodes before proceeding with section [Updating all Controller nodes](#).

Procedure

1. Log into the undercloud as the **stack** user.

2. Source the **stackrc** file:

```
$ source ~/stackrc
```

3. Run the **openstack overcloud external-update run** command against the tasks that have the ovn tag:

```
$ openstack overcloud external-update run --stack <stack_name> --tags ovn
```

- If the name of your overcloud stack is different from the default stack name **overcloud**, set your stack name with the **--stack** option and replace **<stack_name>** with the name of your stack.
4. Wait until the ovn-controller container update completes.

3.4. UPDATING ALL CONTROLLER NODES

Update all the Controller nodes to the latest RHOSP 16.2 version. Run the **openstack overcloud update run** command and include the **--limit Controller** option to restrict operations to the Controller nodes only. The control plane API is fully available during the minor update.



IMPORTANT

Until [BZ#1872404](#) is resolved, for nodes based on composable roles, you must update the **Database** role first, before you can update **Controller, Messaging, Compute, Ceph**, and other roles.



NOTE

If you are not using the default stack name, which is **overcloud**, set your stack name with the **--stack <stack_name>** option and replace **<stack_name>** with the name of your stack.

Procedure

1. Source the **stackrc** file:

```
$ source ~/stackrc
```

2. Run the update command:

```
$ openstack overcloud update run --stack <stack_name> --limit Controller
```

3. Wait until the Controller node update completes.

3.5. UPDATING ALL COMPUTE NODES

Update all Compute nodes to the latest RHOSP 16.2 version. To update Compute nodes, run the **openstack overcloud update run** command and include the **--limit Compute** option to restrict operations to the Compute nodes only.

Parallelization considerations

When you update a large number of Compute nodes, to improve performance, you can run multiple update tasks in the background and configure each task to update a separate group of 20 nodes. For example, if you have 80 Compute nodes in your deployment, you can run the following commands to update the Compute nodes in parallel:

```
$ openstack overcloud update run -y --limit 'Compute[0:19]' > update-compute-0-19.log 2>&1 &
$ openstack overcloud update run -y --limit 'Compute[20:39]' > update-compute-20-39.log 2>&1 &
$ openstack overcloud update run -y --limit 'Compute[40:59]' > update-compute-40-59.log 2>&1 &
$ openstack overcloud update run -y --limit 'Compute[60:79]' > update-compute-60-79.log 2>&1 &
```

This method of partitioning the nodes space is random and you do not have control over which nodes are updated. The selection of nodes is based on the inventory file that you generate when you run the **tripleo-ansible-inventory** command.

To update specific Compute nodes, list the nodes that you want to update in a batch separated by a comma:

```
$ openstack overcloud update run --limit <Compute0>,<Compute1>,<Compute2>,<Compute3>
```



NOTE

If you are not using the default stack name, which is **overcloud**, set your stack name with the **--stack <stack_name>** option and replace **<stack_name>** with the name of your stack.

Procedure

1. Source the **stackrc** file:

```
$ source ~/stackrc
```

2. Run the update command:

```
$ openstack overcloud update run --stack <stack_name> --limit Compute
```

3. Wait until the Compute node update completes.

3.6. UPDATING ALL HCI COMPUTE NODES

Update the Hyperconverged Infrastructure (HCI) Compute nodes to the latest RHOSP 16.2 version. To update the HCI Compute nodes, run the **openstack overcloud update run** command and include the **-limit ComputeHCI** option to restrict operations to only the HCI nodes. You must also run the **openstack overcloud external-update run --tags ceph** command to perform an update to a containerized Red Hat Ceph Storage 4 cluster.

Prerequisites

- On a Ceph Monitor or Controller node that is running the **ceph-mon** service, check that the Red Hat Ceph Storage cluster status is healthy and the pg status is **active+clean**:

```
$ sudo podman exec -it ceph-mon-controller-0 ceph -s
```

If the Ceph cluster is healthy, it returns a status of **HEALTH_OK**.

If the Ceph cluster status is unhealthy, it returns a status of **HEALTH_WARN** or **HEALTH_ERR**. For troubleshooting guidance, see the [Red Hat Ceph Storage 4 Troubleshooting Guide](#) .

Procedure

1. Source the **stackrc** file:

```
$ source ~/stackrc
```

2. Run the update command:

```
$ openstack overcloud update run --stack <stack_name> --limit ComputeHCI
```

- Replace `<stack_name>` with the name of with the name of your stack. if not specified, the default is **overcloud**.
3. Wait until the node update completes.
 4. Run the Ceph Storage update command:

```
$ openstack overcloud external-update run --stack <stack_name> --tags ceph
```

5. Wait until the Compute HCI node update completes.

3.7. UPDATING ALL DISTRIBUTEDCOMPUTEHCI NODES

Update roles specific to distributed compute node architecture. When you upgrade distributed compute nodes, update **DistributedComputeHCI** nodes first, and then update **DistributedComputeHCIScaleOut** nodes.



NOTE

If you are not using the default stack name, which is `overcloud`, set your stack name with the `--stack <stack_name>` option and replace `<_stack_name_>` with the name of your stack.

Prerequisites

- On a Ceph Monitor or Controller node that is running the **ceph-mon** service, check that the Red Hat Ceph Storage cluster status is healthy and the pg status is **active+clean**:

```
$ sudo podman exec -it ceph-mon-controller-0 ceph -s
```

If the Ceph cluster is healthy, it returns a status of **HEALTH_OK**.

If the Ceph cluster status is unhealthy, it returns a status of **HEALTH_WARN** or **HEALTH_ERR**. For troubleshooting guidance, see the [Red Hat Ceph Storage 4 Troubleshooting Guide](#) .

Procedure

1. Source the **stackrc** file:

```
$ source ~/stackrc
```

2. Run the update command:

```
$ openstack overcloud update run --stack <stack_name> --limit DistributedComputeHCI
```

3. Wait until the **DistributedComputeHCI** node update completes.
4. Run the Ceph Storage update command:

```
$ openstack overcloud external-update run --stack <stack_name> --tags ceph
```

5. Wait until the **DistributedComputeHCI** node update completes.

- Use the same process to update **DistributedComputeHCIScaleOut** nodes.

3.8. UPDATING ALL CEPH STORAGE NODES

Update the Red Hat Ceph Storage nodes to the latest RHOSP 16.2 version.



IMPORTANT

RHOSP 16.2 is supported on RHEL 8.4. However, hosts that are mapped to the Ceph Storage role update to the latest major RHEL release. For more information, see [Red Hat Ceph Storage: Supported configurations](#).

Prerequisites

- On a Ceph Monitor or Controller node that is running the **ceph-mon** service, check that the Red Hat Ceph Storage cluster status is healthy and the pg status is **active+clean**:

```
$ sudo podman exec -it ceph-mon-controller-0 ceph -s
```

If the Ceph cluster is healthy, it returns a status of **HEALTH_OK**.

If the Ceph cluster status is unhealthy, it returns a status of **HEALTH_WARN** or **HEALTH_ERR**. For troubleshooting guidance, see the [Red Hat Ceph Storage 4 Troubleshooting Guide](#).

Procedure

- Source the **stackrc** file:

```
$ source ~/stackrc
```

- Update group nodes.

To update all nodes in a group:

```
$ openstack overcloud update run --limit <GROUP_NAME>
```

To update a single node in a group:

```
$ openstack overcloud update run --limit <GROUP_NAME> [NODE_INDEX]
```



NOTE

Ensure that you update all nodes if you choose to update nodes individually.

The index of the first node in a group is zero (0). For example, to update the first node in a group named **CephStorage**:

```
openstack overcloud update run --limit CephStorage[0]
```

- Wait until the node update completes.
- Run the Ceph Storage container update command to run **ceph-ansible** as an external process and update the Red Hat Ceph Storage 4 containers:

-

```
$ openstack overcloud external-update run --tags ceph
```

5. Wait until the Ceph Storage container update completes.

3.9. PERFORMING ONLINE DATABASE UPDATES

Some overcloud components require an online update or migration of their databases tables. To perform online database updates, run the **openstack overcloud external-update run** command against tasks that have the **online_upgrade** tag.

Online database updates apply to the following components:

- OpenStack Block Storage (cinder)
- OpenStack Compute (nova)

Procedure

1. Source the **stackrc** file:

```
$ source ~/stackrc
```

2. Run the **openstack overcloud external-update run** command against tasks that use the **online_upgrade** tag:

```
$ openstack overcloud external-update run --tags online_upgrade
```

3.10. FINALIZING THE UPDATE



NOTE

You are no longer required to run the **openstack overcloud update converge** command. However, if you disabled fencing and plan to skip the converge step, you must manually re-enable fencing.

You can update the overcloud stack to the latest RHOSP 16.2 version. This ensures that the stack resource structure aligns with a standard deployment of OSP 16.2 and you can perform regular **openstack overcloud deploy** functions in the future.

Procedure

1. Log in to the undercloud as the **stack** user.
2. Source the **stackrc** file:

```
$ source ~/stackrc
```

3. If fencing is disabled, and you do not run **openstack overcloud update converge**, then you must re-enable fencing:
 - a. Log in to a Controller node and run the Pacemaker command to re-enable fencing:


```
$ ssh tripleo-admin@<controller_ip> "sudo pcs property set stonith-enabled=true"
```

- Replace **<controller_ip>** with the IP address of a Controller node. You can find the IP addresses of your Controller nodes with the **openstack server list** command.

b. In the **fencing.yaml** environment file, set the value of the **EnableFencing** parameter to **true**.

4. Optional: Run the update finalization command:

```
$ openstack overcloud update converge \
  --templates \
  --stack <stack_name> \
  -r <roles_data_file> \
  -n <network_data_file> \
  -e <environment_file> \
  -e <environment_file> \
  ...
  ...
```

Include the following options that are relevant to your environment:

- The **fencing.yaml** environment file, with the **EnableFencing** parameter set to **true**.
 - If the name of your overcloud stack is different to the default name **overcloud**, include the **--stack** option in the update preparation command and replace **<stack_name>** with the name of your stack.
 - If you use custom roles, include your custom roles (**<roles_data>**) file (**-r**).
 - If you use custom networks, include your composable network (**<network_data>**) file (**-n**).
 - Any custom configuration environment files (**-e**).
- Wait until the update finalization completes.

CHAPTER 4. REBOOTING THE OVERCLOUD

After you perform a minor Red Hat OpenStack Platform (RHOSP) update to the latest 16.2 version, reboot your overcloud. The reboot refreshes the nodes with any associated kernel, system-level, and container component updates. These updates provide performance and security benefit. Plan downtime to perform the reboot procedures.

Use the following guidance to understand how to reboot different node types:

- If you reboot all nodes in one role, reboot each node individually. If you reboot all nodes in a role simultaneously, service downtime can occur during the reboot operation.
- Complete the reboot procedures on the nodes in the following order:
 1. [Section 4.1, “Rebooting Controller and composable nodes”](#)
 2. [Section 4.2, “Rebooting a Ceph Storage \(OSD\) cluster”](#)
 3. [Section 4.3, “Rebooting Compute nodes”](#)

4.1. REBOOTING CONTROLLER AND COMPOSABLE NODES

Reboot Controller nodes and standalone nodes based on composable roles, and exclude Compute nodes and Ceph Storage nodes.

Procedure

1. Log in to the node that you want to reboot.
2. Optional: If the node uses Pacemaker resources, stop the cluster:

```
[heat-admin@overcloud-controller-0 ~]$ sudo pcs cluster stop
```

3. Reboot the node:

```
[heat-admin@overcloud-controller-0 ~]$ sudo reboot
```

4. Wait until the node boots.

Verification

1. Verify that the services are enabled.
 - a. If the node uses Pacemaker services, check that the node has rejoined the cluster:

```
[heat-admin@overcloud-controller-0 ~]$ sudo pcs status
```

- b. If the node uses Systemd services, check that all services are enabled:

```
[heat-admin@overcloud-controller-0 ~]$ sudo systemctl status
```

- c. If the node uses containerized services, check that all containers on the node are active:

```
[heat-admin@overcloud-controller-0 ~]$ sudo podman ps
```

-

4.2. REBOOTING A CEPH STORAGE (OSD) CLUSTER

Complete the following steps to reboot a cluster of Ceph Storage (OSD) nodes.

Prerequisites

- On a Ceph Monitor or Controller node that is running the **ceph-mon** service, check that the Red Hat Ceph Storage cluster status is healthy and the pg status is **active+clean**:

```
$ sudo podman exec -it ceph-mon-controller-0 ceph -s
```


If the Ceph cluster is healthy, it returns a status of **HEALTH_OK**.

If the Ceph cluster status is unhealthy, it returns a status of **HEALTH_WARN** or **HEALTH_ERR**. For troubleshooting guidance, see the [Red Hat Ceph Storage 4 Troubleshooting Guide](#) .

Procedure

- Log in to a Ceph Monitor or Controller node that is running the **ceph-mon** service, and disable Ceph Storage cluster rebalancing temporarily:

```
$ sudo podman exec -it ceph-mon-controller-0 ceph osd set noout
$ sudo podman exec -it ceph-mon-controller-0 ceph osd set norebalance
```



NOTE

If you have a multistack or distributed compute node (DCN) architecture, you must specify the cluster name when you set the **noout** and **norebalance** flags. For example: **sudo podman exec -it ceph-mon-controller-0 ceph osd set noout --cluster <cluster_name>**

- Select the first Ceph Storage node that you want to reboot and log in to the node.
- Reboot the node:

```
$ sudo reboot
```


- Wait until the node boots.
- Log into the node and check the cluster status:

```
$ sudo podman exec -it ceph-mon-controller-0 ceph status
```

Check that the **pgmap** reports all **pgs** as normal (**active+clean**).

- Log out of the node, reboot the next node, and check its status. Repeat this process until you have rebooted all Ceph Storage nodes.
- When complete, log in to a Ceph Monitor or Controller node that is running the **ceph-mon** service and re-enable cluster rebalancing:

```
$ sudo podman exec -it ceph-mon-controller-0 ceph osd unset noout
$ sudo podman exec -it ceph-mon-controller-0 ceph osd unset norebalance
```



NOTE

If you have a multistack or distributed compute node (DCN) architecture, you must specify the cluster name when you unset the **noout** and **norebalance** flags. For example: **sudo podman exec -it ceph-mon-controller-0 ceph osd set noout --cluster <cluster_name>**

8. Perform a final status check to verify that the cluster reports **HEALTH_OK**:

```
$ sudo podman exec -it ceph-mon-controller-0 ceph status
```

4.3. REBOOTING COMPUTE NODES

To ensure minimal downtime of instances in your Red Hat OpenStack Platform environment, the [Migrating instances workflow](#) outlines the steps you must complete to migrate instances from the Compute node that you want to reboot.



NOTE

If you do not migrate the instances from the source Compute node to another Compute node, the instances might be restarted on the source Compute node, which might cause the upgrade to fail. This is related to the known issue around changes to Podman and the libvirt service:

- [BZ#2009106 - podman panic after tripleo_nova_libvirt restart two times](#)
- [BZ#2010135 - podman panic after tripleo_nova_libvirt restart two times](#)

Migrating instances workflow

1. Decide whether to migrate instances to another Compute node before rebooting the node.
2. Select and disable the Compute node that you want to reboot so that it does not provision new instances.
3. Migrate the instances to another Compute node.
4. Reboot the empty Compute node.
5. Enable the empty Compute node.

Prerequisites

- Before you reboot the Compute node, you must decide whether to migrate instances to another Compute node while the node is rebooting. Review the list of migration constraints that you might encounter when you migrate virtual machine instances between Compute nodes. For more information, see [Migration constraints](#) in *Configuring the Compute Service for Instance Creation*.

- If you cannot migrate the instances, you can set the following core template parameters to control the state of the instances after the Compute node reboots:

NovaResumeGuestsStateOnHostBoot

Determines whether to return instances to the same state on the Compute node after reboot. When set to **False**, the instances remain down and you must start them manually. The default value is **False**.

NovaResumeGuestsShutdownTimeout

Number of seconds to wait for an instance to shut down before rebooting. It is not recommended to set this value to **0**. The default value is **300**. For more information about overcloud parameters and their usage, see [Overcloud Parameters](#).

Procedure

1. Log in to the undercloud as the **stack** user.
2. List all Compute nodes and their UUIDs:

```
$ source ~/stackrc
(undercloud) $ openstack server list --name compute
```

Identify the UUID of the Compute node that you want to reboot.

3. From the undercloud, select a Compute node and disable it:

```
$ source ~/overcloudrc
(overcloud) $ openstack compute service list
(overcloud) $ openstack compute service set <hostname> nova-compute --disable
```

4. List all instances on the Compute node:

```
(overcloud) $ openstack server list --host <hostname> --all-projects
```

5. Optional: If you decide to migrate the instances to another Compute node, complete the following steps:

- a. If you decide to migrate the instances to another Compute node, use one of the following commands:

- To migrate the instance to a different host, run the following command:

```
(overcloud) $ openstack server migrate <instance_id> --live <target_host> --wait
```

- Let **nova-scheduler** automatically select the target host:

```
(overcloud) $ nova live-migration <instance_id>
```

- Live migrate all instances at once:

```
$ nova host-evacuate-live <hostname>
```

**NOTE**

The **nova** command might cause some deprecation warnings, which are safe to ignore.

- b. Wait until migration completes.
 - c. Confirm that the migration was successful:

```
(overcloud) $ openstack server list --host <hostname> --all-projects
```
 - d. Continue to migrate instances until none remain on the Compute node.
6. Log in to the Compute node and reboot the node:

```
[heat-admin@overcloud-compute-0 ~]$ sudo reboot
```
 7. Wait until the node boots.
 8. Re-enable the Compute node:

```
$ source ~/overcloudrc  
(overcloud) $ openstack compute service set <hostname> nova-compute --enable
```
 9. Check that the Compute node is enabled:

```
(overcloud) $ openstack compute service list
```