



Red Hat OpenStack Platform 17.1

Performing a minor update of Red Hat OpenStack Platform

Apply the latest bug fixes and security improvements to Red Hat OpenStack Platform

Red Hat OpenStack Platform 17.1 Performing a minor update of Red Hat OpenStack Platform

Apply the latest bug fixes and security improvements to 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	6
Procedure	6
1.1. UPGRADE PATHS FOR LONG LIFE RELEASES	7
1.2. LOCKING THE ENVIRONMENT TO A RED HAT ENTERPRISE LINUX RELEASE	7
1.3. UPDATING RED HAT OPENSTACK PLATFORM REPOSITORIES	8
1.4. UPDATING THE CONTAINER IMAGE PREPARATION FILE	10
1.5. DISABLING FENCING IN THE OVERCLOUD	10
CHAPTER 2. UPDATING THE UNDERCLOUD	12
Prerequisites	12
2.1. VALIDATING RHOSP BEFORE THE UNDERCLOUD UPDATE	12
2.2. VALIDATING YOUR SSH KEY SIZE	13
2.3. PERFORMING A MINOR UPDATE OF A CONTAINERIZED UNDERCLOUD	13
2.4. UPDATING THE OVERCLOUD IMAGES	14
CHAPTER 3. UPDATING THE OVERCLOUD	16
Prerequisites	16
Procedure	16
3.1. RUNNING THE OVERCLOUD UPDATE PREPARATION	16
3.2. RUNNING THE CONTAINER IMAGE PREPARATION	18
3.3. OPTIONAL: UPDATING THE OVN-CONTROLLER CONTAINER ON ALL OVERCLOUD SERVERS	18
3.4. UPDATING THE CONTAINER IMAGE NAMES OF PACEMAKER-CONTROLLED SERVICES	19
3.5. UPDATING ALL CONTROLLER NODES	19
3.6. UPDATING ALL COMPUTE NODES	20
3.7. UPDATING ALL HCI COMPUTE NODES	21
3.8. UPDATING ALL DISTRIBUTEDCOMPUTEHCI NODES	21
3.9. UPDATING ALL CEPH STORAGE NODES	22
3.10. UPDATING THE RED HAT CEPH STORAGE CLUSTER	23
3.11. PERFORMING ONLINE DATABASE UPDATES	24
3.12. RE-ENABLING FENCING IN THE OVERCLOUD	24
CHAPTER 4. REBOOTING THE OVERCLOUD	26
4.1. REBOOTING CONTROLLER AND COMPOSABLE NODES	26
4.2. REBOOTING A CEPH STORAGE (OSD) CLUSTER	27
4.3. REBOOTING COMPUTE NODES	28
4.4. VALIDATING RHOSP AFTER THE OVERCLOUD UPDATE	30

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) 17.1 environment updated with the latest packages and containers.

Use the upgrade path for the following versions:


Old RHOSP Version	New RHOSP Version
Red Hat OpenStack Platform 17.0.z	Red Hat OpenStack Platform 17.1 latest
Red Hat OpenStack Platform 17.1.z	Red Hat OpenStack Platform 17.1 latest

Minor update workflow

A minor update of your RHOSP environment involves updating the RPM packages and containers on the undercloud and overcloud host, and the service configuration, if needed. The data plane and control plane are fully available during the minor update. You must complete each of the following steps to update your RHOSP environment:

Update step	Description
Undercloud update	Director packages are updated, containers are replaced, and the undercloud is rebooted.
Optional ovn-controller update	All ovn-controller containers are updated in parallel on all Compute and Controller hosts.
ha-image-update external	Updates container image names of Pacemaker-controlled services. There is no service disruption. This step applies to only customers that are updating their system from version 17.0.z to the latest 17.1 release.
Overcloud update of Controller nodes and composable nodes that contain Pacemaker services	During an Overcloud update, the Pacemaker services are stopped for each host. While the Pacemaker services are stopped, the RPMs on the host, the container configuration data, and the containers are updated. When the Pacemaker services restart, the host is added again.
Overcloud update of composable nodes without Pacemaker services	Networker, ObjectStorage, BlockStorage, or any other role that does not include Pacemaker services are updated one node at a time.
Overcloud update of Compute nodes	Multiple nodes are updated in parallel. The default value for running nodes in parallel is 25.
Overcloud update of Ceph nodes	Ceph nodes are updated one node at a time.

Update step	Description
Ceph cluster update	Ceph services are updated by using cephadm . The update occurs per daemon, beginning with CephMgr , CephMon , CephOSD , and then additional daemons.



NOTE

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.

Additionally, an administrator can perform the following operations during a minor update:

- Migrate your virtual machine
- Create a virtual machine network
- Run additional cloud operations

The following operations are not supported during a minor update:

- Replacing a Controller node
- Scaling in or scaling out any role

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

There are currently no known issues.

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, "Updating Red Hat Openstack Platform repositories"](#)

3. [Section 1.4, “Updating the container image preparation file”](#)
4. [Section 1.5, “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.



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 17.0.x on RHEL 9.0	RHOSP 17.0 latest on RHEL 9.0 latest
RHOSP 17.1.x on RHEL 9.2	RHOSP 17.1 latest on RHEL 9.2 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
RHOSP 16 on RHEL 8.4	RHOSP 17.1 latest on RHEL 9.0 latest

For more information, see [Framework for upgrades \(16.2 to 17.1\)](#).

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

Red Hat OpenStack Platform (RHOSP) 17.1 is supported on Red Hat Enterprise Linux (RHEL) 9.2. Before you perform the update, lock the undercloud and overcloud repositories to the RHEL 9.2 release to avoid upgrading the operating system to a newer minor release.

Procedure

1. Log in to the undercloud host as the **stack** user.
2. Source the **stackrc** undercloud credentials 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 9.2:

```
parameter_defaults:
  RhsmVars:
    ...
    rhsm_username: "myusername"
    rhsm_password: "p@55w0rd!"
    rhsm_org_id: "1234567"
    rhsm_pool_ids: "1a85f9223e3d5e43013e3d6e8ff506fd"
    rhsm_method: "portal"
    rhsm_release: "9.2"
```

5. Save the overcloud subscription management environment file.
6. Create a playbook that contains a task to lock the operating system version to RHEL 9.2 on all nodes:

```
$ cat > ~/set_release.yml <<'EOF'
- hosts: all
  gather_facts: false
  tasks:
    - name: set release to 9.2
      command: subscription-manager release --set=9.2
      become: true
EOF
```

7. Run the **set_release.yml** playbook:

```
$ ansible-playbook -i ~/overcloud-deploy/<stack>/tripleo-ansible-inventory.yml -f 25
~/set_release.yml --limit <undercloud>, <Controller>, <Compute>
```

- Replace **<stack>** with the name of your stack.
- 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. Do not run this playbook against Ceph Storage nodes because you might have 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=9.2
```

1.3. UPDATING RED HAT OPENSTACK PLATFORM REPOSITORIES

Update your repositories to use Red Hat OpenStack Platform (RHOSP) 17.1.

Procedure

1. Log in to the undercloud host as the **stack** user.
2. Source the **stackrc** undercloud credentials 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 is using the RHOSP 17.1 repositories, change the repository to the correct versions:

```
parameter_defaults:
  RhsmVars:
    rhsm_repos:
      - rhel-9-for-x86_64-baseos-eus-rpms
      - rhel-9-for-x86_64-appstream-eus-rpms
      - rhel-9-for-x86_64-highavailability-eus-rpms
      - openstack-17.1-for-rhel-9-x86_64-rpms
      - fast-datapath-for-rhel-9-x86_64-rpms
```

5. Save the overcloud subscription management environment file.
6. Create a playbook that contains a task to set the repositories to RHOSP 17.1 on all nodes:

```
$ cat > ~/update_rhosp_repos.yaml <<'EOF'
- hosts: all
  gather_facts: false
  tasks:
    - name: change osp repos
      command: subscription-manager repos --enable=openstack-17.1-for-rhel-9-x86_64-rpms
      become: true
EOF
```

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

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

- Replace **<stack>** with the name of your stack.
 - Use the **--limit** option to apply the content to all RHOSP nodes. Replace **<undercloud>**, **<Controller>**, and **<Compute>** with the Ansible groups in your environment that contain those nodes. Do not run this playbook against Ceph Storage nodes because they usually use a different subscription.
8. Create a playbook that contains a task to set the repositories to RHOSP 17.1 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 --enable=openstack-17.1-deployment-tools-for-
rhel-9-x86_64-rpms
  become: true
EOF

```

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

```

$ ansible-playbook -i ~/overcloud-deploy/<stack>/tripleo-ansible-inventory.yaml -f 25
~/update_ceph_repos.yaml --limit CephStorage

```

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

1.4. 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

- Edit the container preparation file. The default name for this file is usually **containers-prepare-parameter.yaml**.
- Ensure that the **tag** parameter is set to **17.1** for each rule set:

```

parameter_defaults:
  ContainerImagePrepare:
    - push_destination: true
      set:
        ...
        tag: '17.1'
        tag_from_label: '{version}-{release}'

```



NOTE

If you do not want to use a specific tag for the update, such as **17.1** or **17.1.1**, 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.

- Save this file.

1.5. 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.

Procedure

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

```
$ source ~/stackrc
```

3. For each Controller node, log in to the Controller node and run the Pacemaker command to disable fencing:

```
$ ssh tripleo-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 at **/etc/hosts** or **/var/lib/mistral**.
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) 17.1 version, complete the following procedures:

1. [Section 2.3, “Performing a minor update of a containerized undercloud”](#)
2. [Section 2.4, “Updating the overcloud images”](#)

Prerequisites

- Before you can update the undercloud to the latest RHOSP 17.1 version, ensure that you complete all the update preparation procedures. For more information, see [Chapter 1, *Preparing for a minor update*](#)

2.1. VALIDATING RHOSP BEFORE THE UNDERCLOUD UPDATE

Before you update your Red Hat OpenStack Platform (RHOSP) environment, validate your undercloud with the **tripleo-validations** playbooks.

For more information about validations, see [Using the validation framework](#) in *Installing and managing Red Hat OpenStack Platform with director*.

Procedure

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

```
$ source ~/stackrc
```

3. Install the packages for validation:

```
$ sudo dnf -y update openstack-tripleo-validations python3-validations-libs validations-common
```

4. Run the validation:

```
$ validation run -i ~/overcloud-deploy/<stack>/tripleo-ansible-inventory.yaml --group pre-update
```

- Replace `<stack>` with the name of the stack.

Verification

1. To view the results of the validation report, see [Viewing validation history](#) in *Installing and managing Red Hat OpenStack Platform with director*.



NOTE

If a host is not found when you run a validation, the command reports the status as **SKIPPED**. A status of **SKIPPED** means that the validation is not executed, which is expected. Additionally, if a validation's pass criteria is not met, the command reports the status as **FAILED**. A **FAILED** validation does not prevent you from using your updated RHOSP environment. However, a **FAILED** validation can indicate an issue with your environment.

2.2. VALIDATING YOUR SSH KEY SIZE

Starting with Red Hat Enterprise Linux (RHEL) 9.1, a minimum SSH key size of 2048 bits is required. If your current SSH key on Red Hat OpenStack Platform (RHOSP) director is less than 2048 bits, you can lose access to the overcloud. You must verify that your SSH key meets the required bit size.

Procedure

1. Validate your SSH key size:

```
ssh-keygen -l -f ~/.ssh/id_rsa.pub
```

Example output:

```
1024 SHA256:Xqz0Xz0/aJua6B3qRD7VsLr6n/V3zhmnGskcFR6FIJw
stack@director.example.local (RSA)
```

2. If your SSH key is less than 2048 bits, you must rotate out the SSH key before continuing. For more information, see [Updating SSH keys in your OpenStack environment](#) in *Hardening Red Hat OpenStack Platform*.

2.3. 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. Log in to the undercloud host as the **stack** user.
2. Source the **stackrc** undercloud credentials file:

```
$ source ~/stackrc
```

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

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

4. Update the undercloud environment:

```
$ 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.

2.4. 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.

Prerequisites

- You have updated the undercloud node to the latest version. For more information, see [Section 2.3, "Performing a minor update of a containerized undercloud"](#).

Procedure

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

```
$ source ~/stackrc
```

3. Remove any existing images from the **images** directory on the **stack** user's home (**/home/stack/images**):

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

4. Extract the archives:

```
cd ~/images
for i in /usr/share/rhosp-director-images/ironic-python-agent-latest-17.1.tar /usr/share/rhosp-
director-images/overcloud-hardened-uefi-full-latest-17.1.tar; do tar -xvf $i; done
cd ~
```

5. Import the latest images into the director:

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

6. Configure your nodes to use the new images:

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

7. Verify the existence of the new images:

```
$ ls -l /var/lib/ironic/httpboot /var/lib/ironic/images
```



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 17.1 images with the RHOSP 17.1 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, and updating your nodes. 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 [Understanding heat templates](#) in the *Customizing your Red Hat OpenStack Platform deployment* guide. You must update the local copy before you upgrade the overcloud.
- Add the **GlanceApiInternal** service to your Controller role:

```
OS::TripleO::Services::GlanceApiInternal
```

This is the service for the internal instance of the Image service (glance) API to provide location data to administrators and other services that require it, such as the Block Storage service (cinder) and the Compute service (nova).

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 the container image names of Pacemaker-controlled services"](#)
5. [Section 3.5, "Updating all Controller nodes"](#)
6. [Section 3.6, "Updating all Compute nodes"](#)
7. [Section 3.7, "Updating all HCI Compute nodes"](#)
8. [Section 3.8, "Updating all DistributedComputeHCI nodes"](#)
9. [Section 3.9, "Updating all Ceph Storage nodes"](#)
10. [Section 3.10, "Updating the Red Hat Ceph Storage cluster"](#)
11. [Section 3.11, "Performing online database updates"](#)
12. [Section 3.12, "Re-enabling fencing in the overcloud"](#)

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) 17.1 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 [Defining custom network interface templates](#) in the *Installing and managing Red Hat OpenStack Platform with director* 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. Log in to the undercloud host as the **stack** user.
2. Source the **stackrc** undercloud credentials file:

```
$ source ~/stackrc
```

3. 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, use the **-r** option to include the custom roles (**<roles_data_file>**) file.
- If you use custom networks, use the **-n** option to include your composable network in the (**<network_data_file>**) file.
- 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.
- Include any custom configuration environment files with the **-e** option.

4. 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 17.1 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.

Procedure

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

```
$ source ~/stackrc
```

3. 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
```

- 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.

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 17.1 version. The update occurs on every overcloud server that runs the **ovn-controller** container.

- 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.
- 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](#).

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

Procedure

1. Log in to the undercloud host as the **stack** user.
2. Source the **stackrc** undercloud credentials 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 THE CONTAINER IMAGE NAMES OF PACEMAKER-CONTROLLED SERVICES

If you update your system from Red Hat Openstack Platform (RHOSP) 17 to RHOSP 17.1, you must update the container image names of the Pacemaker-controlled services. You must perform this update to migrate to the new image naming schema of the Pacemaker-controlled services.

If you update your system from a version of RHOSP 17.1 to the latest version of RHOSP 17.1, you do not need to update the container image names of the pacemaker-controlled services.

Procedure

1. Log in to the undercloud host as the stack user.
2. Source the `stackrc` undercloud credentials file:

```
$ source ~/stackrc
```

3. Run the **openstack overcloud external-update run** command with the **ha_image_update** tag:

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

- If the name of your undercloud stack is different from the default stack name `undercloud`, set your stack name with the **--stack** option and replace `<stack_name>` with the name of your stack.

3.5. UPDATING ALL CONTROLLER NODES

Update all the Controller nodes to the latest RHOSP 17.1 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.

Procedure

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

```
$ source ~/stackrc
```

3. Run the update command:

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

- 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 Controller node update completes.

3.6. UPDATING ALL COMPUTE NODES

Update all Compute nodes to the latest RHOSP 17.1 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>
```

Procedure

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

```
$ source ~/stackrc
```

3. Run the update command:

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

- 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 Compute node update completes.

3.7. UPDATING ALL HCI COMPUTE NODES

Update the Hyperconverged Infrastructure (HCI) Compute nodes to the latest RHOSP 17.1 version.

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 cephadm -- shell ceph status
```

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 5 Troubleshooting Guide](#) or the [Red Hat Ceph Storage 6 Troubleshooting Guide](#)

Procedure

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

```
$ source ~/stackrc
```

3. Run the update command:

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

- 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 node update completes.

3.8. 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.

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 cephadm -- shell ceph status
```

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 5 Troubleshooting Guide](#) or [Red Hat Ceph Storage 6 Troubleshooting Guide](#).

Procedure

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

```
$ source ~/stackrc
```


3. Run the update command:

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

- 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 **DistributedComputeHCI** node update completes.
 5. Use the same process to update **DistributedComputeHCIScaleOut** nodes.

3.9. UPDATING ALL CEPH STORAGE NODES

Update the Red Hat Ceph Storage nodes to the latest RHOSP 17.1 version.



IMPORTANT

RHOSP 17.1 is supported on RHEL 9.2. 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 cephadm -- shell ceph status
```

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 5 Troubleshooting Guide](#) or the [Red Hat Ceph Storage 6 Troubleshooting Guide](#).

Procedure

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

```
$ source ~/stackrc
```

3. Run the update command:

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

-

- 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 node update completes.

3.10. UPDATING THE RED HAT CEPH STORAGE CLUSTER

Update the director-deployed Red Hat Ceph Storage cluster to the latest version that is compatible with Red Hat OpenStack Platform (RHOSP) 17.1 by using the **cephadm** command.

Update your Red Hat Ceph Storage cluster if one of the following scenarios applies to your environment:

- If you upgraded from RHOSP 16.2 to RHOSP 17.1, you run Red Hat Ceph Storage 5, and you are updating to a newer version of Red Hat Ceph Storage 5.
- If you newly deployed RHOSP 17.1, you run Red Hat Ceph Storage 6, and you are updating to a newer version of Red Hat Ceph Storage 6.

Prerequisites

- Complete the container image preparation in [Section 3.2, "Running the container image preparation"](#).

Procedure

1. Log in to a Controller node.
2. Check the health of the cluster:

```
$ sudo cephadm shell -- ceph health
```



NOTE

If the Ceph Storage cluster is healthy, the command returns a result of **HEALTH_OK**. If the command returns a different result, review the status of the cluster and contact Red Hat support before continuing the update. For more information, see [Upgrade a Red Hat Ceph Storage cluster using cephadm](#) in the *Red Hat Ceph Storage Upgrade Guide* or [Upgrade a Red Hat Ceph Storage cluster using cephadm](#) in the *Red Hat Ceph Storage 6 Upgrade Guide*.

3. Optional: Check which images should be included in the Ceph Storage cluster update:

```
$ openstack tripleo container image list -f value | awk -F '/' '{print $2}'
```

4. Update the cluster to the latest Red Hat Ceph Storage version:

```
$ sudo cephadm shell -- ceph orch upgrade start --image <image_name>: <version>
```

- Replace **<image_name>** with the name of the Ceph Storage cluster image.

- Replace **<version>** with the target version to which you are updating the Ceph Storage cluster.
5. Wait until the Ceph Storage container update completes. To monitor the status of the update, run the following command:

```
sudo cephadm shell -- ceph orch upgrade status
```

3.11. 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. Log in to the undercloud host as the **stack** user.
2. Source the **stackrc** undercloud credentials file:

```
$ source ~/stackrc
```

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

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

3.12. RE-ENABLING FENCING IN THE OVERCLOUD

Before you updated the overcloud, you disabled fencing in [Disabling fencing in the overcloud](#). After you update the overcloud, re-enable fencing to protect your data if a node fails.

Procedure

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

```
$ source ~/stackrc
```

3. 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.

4. In the **fencing.yaml** environment file, set the **EnableFencing** parameter to **true**.

Additional Resources

- [Fencing Controller nodes with STONITH](#)

CHAPTER 4. REBOOTING THE OVERCLOUD

After you perform a minor Red Hat OpenStack Platform (RHOSP) update to the latest 17.1 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 benefits. 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:

```
[tripleo-admin@overcloud-controller-0 ~]$ sudo pcs cluster stop
```
3. Reboot the node:

```
[tripleo-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:

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

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

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

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

```
[tripleo-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 cephadm -- shell ceph status
```

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 5 Troubleshooting Guide](#) or the [Red Hat Ceph Storage 6 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 cephadm shell -- ceph osd set noout
$ sudo cephadm shell -- ceph osd set norebalance
```



NOTE

If you have a multistack or distributed compute node (DCN) architecture, you must specify the Ceph cluster name when you set the **noout** and **norebalance** flags. For example: **sudo cephadm shell -c /etc/ceph/<cluster>.conf -k /etc/ceph/<cluster>.client.keyring**.

- 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 in to the node and check the Ceph cluster status:

```
$ sudo cephadm -- shell 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 enable Ceph cluster rebalancing:

```
$ sudo cephadm shell -- ceph osd unset noout
$ sudo cephadm shell -- ceph osd unset norebalance
```



NOTE

If you have a multistack or distributed compute node (DCN) architecture, you must specify the Ceph cluster name when you unset the **noout** and **norebalance** flags. For example: **sudo cephadm shell -c /etc/ceph/<cluster>.conf -k /etc/ceph/<cluster>.client.keyring**

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

```
$ sudo cephadm shell 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.

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*.



NOTE

If you have a Multi-RHEL environment, and you want to migrate virtual machines from a Compute node that is running RHEL 9.2 to a Compute node that is running RHEL 8.4, only cold migration is supported. For more information about cold migration, see [Cold migrating an instance](#) 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. Retrieve a list of your Compute nodes to identify the host name of the node that you want to reboot:

```
(undercloud)$ source ~/overcloudrc
(overcloud)$ openstack compute service list
```

Identify the host name of the Compute node that you want to reboot.

3. Disable the Compute service on the Compute node that you want to reboot:

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

- Replace **<hostname>** with the host name of your Compute node.

4. List all instances on the Compute node:

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

5. Optional: 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
```

- Replace **<instance_id>** with your instance ID.
- Replace **<target_host>** with the host that you are migrating the instance to.

- 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:

```
[tripleo-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
```

4.4. VALIDATING RHOSP AFTER THE OVERCLOUD UPDATE

After you update your Red Hat OpenStack Platform (RHOSP) environment, validate your overcloud with the **tripleo-validations** playbooks.

For more information about validations, see [Using the validation framework](#) in *Installing and managing Red Hat OpenStack Platform with director*.

Procedure

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

```
$ source ~/stackrc
```


3. Run the validation:

```
$ validation run -i ~/overcloud-deploy/<stack>/tripleo-ansible-inventory.yaml --group post-update
```

- Replace <stack> with the name of the stack.

Verification

1. To view the results of the validation report, see [Viewing validation history](#) in *Installing and managing Red Hat OpenStack Platform with director*.



NOTE

If a host is not found when you run a validation, the command reports the status as **SKIPPED**. A status of **SKIPPED** means that the validation is not executed, which is expected. Additionally, if a validation's pass criteria is not met, the command reports the status as **FAILED**. A **FAILED** validation does not prevent you from using your updated RHOSP environment. However, a **FAILED** validation can indicate an issue with your environment.