Red Hat OpenStack Platform 15

Keeping Red Hat OpenStack Platform Updated

Performing minor updates of Red Hat OpenStack Platform

Red Hat OpenStack Platform 15 Keeping Red Hat OpenStack Platform Updated

Performing minor updates of Red Hat OpenStack Platform

OpenStack Team
rhos-docs@redhat.com

Legal Notice

Abstract

This document provides the procedure to update your Red Hat OpenStack Platform 15 (Stein) environment. This document assumes you will update a containerized OpenStack Platform deployment installed on Red Hat Enterprise Linux 8.

Table of Contents

CHAPTER 1. INTRODUCTION

This document provides a workflow to help keep your Red Hat OpenStack Platform 15 environment updated with the latest packages and containers.

This guide provides an upgrade path through the following versions:

Old Overcloud Version	New Overcloud Version
Red Hat OpenStack Platform 15	Red Hat OpenStack Platform 15.z

1.1. HIGH LEVEL WORKFLOW

The following table provides an outline of the steps required for the upgrade process:

Step	Description
Updating the undercloud	Update the undercloud to the latest OpenStack Platform 15.z version.
Updating the overcloud	Update the overcloud to the latest OpenStack Platform 15.z version.
Updating the Ceph Storage nodes	Upgrade all Ceph Storage services.
Finalize the upgrade	Run the convergence command to refresh your overcloud stack.

CHAPTER 2. UPDATING THE UNDERCLOUD

This process updates the undercloud and its overcloud images to the latest **Red Hat OpenStack Platform 15** version.

2.1. PERFORMING A MINOR UPDATE OF A CONTAINERIZED UNDERCLOUD

The director provides commands to update the packages on the undercloud node. This allows you to perform a minor update within the current version of your OpenStack Platform environment.

Procedure

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

2. Run **dnf** to upgrade the director's main packages:

   ```
   $ sudo dnf update -y python3-tripleoclient* openstack-tripleo-common openstack-tripleo-heat-templates
   ```

3. The director uses the **openstack undercloud upgrade** command to update the undercloud environment. Run the command:

   ```
   $ openstack undercloud upgrade
   ```

4. Wait until the undercloud upgrade process completes.

5. Reboot the undercloud to update the operating system's kernel and other system packages:

   ```
   $ sudo reboot
   ```

6. Wait until the node boots.

2.2. UPDATING THE OVERCLOUD IMAGES

You need to replace your current overcloud images with new versions. The new images ensure the director can introspect and provision your nodes using the latest version of OpenStack Platform software.

Prerequisites

- You have updated the undercloud to the latest version.

Procedure

1. Source the **stackrc** file:

   ```
   $ source ~/stackrc
   ```

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

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

3. Extract the archives:

```
$ cd ~/images
$ for i in /usr/share/rhosp-director-images/overcloud-full-latest-15.0.tar /usr/share/rhosp-director-images/ironic-python-agent-latest-15.0.tar; do tar -xvf $i; done
$ cd ~
```

4. Import the latest images into the director:

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

5. Configure your nodes to use the new images:

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

6. Verify the existence of the new images:

```
$ openstack image list
$ ls -l /httpboot
```

> **IMPORTANT**
>
> When deploying overcloud nodes, ensure the Overcloud image version corresponds to the respective Heat template version. For example, only use the OpenStack Platform 15 images with the OpenStack Platform 15 Heat templates.

2.3. UNDERCLOUD POST-UPGRADE NOTES

- If using a local set of core templates in your **stack** users home directory, ensure you update the templates using the recommended workflow in Using Customized Core Heat Templates. You must update the local copy before upgrading the overcloud.

2.4. NEXT STEPS

The undercloud upgrade is complete. You can now update the overcloud.

CHAPTER 3. UPDATING THE OVERCLOUD

This process updates the overcloud.

Prerequisites

- You have updated the undercloud to the latest version.

3.1. RUNNING THE OVERCLOUD UPDATE PREPARATION

The update requires running **openstack overcloud update prepare** command, which performs the following tasks:

- Updates the overcloud plan to OpenStack Platform 15

- Prepares the nodes for the update

Procedure

1. Source the **stackrc** file:

   ```
   $ source ~/stackrc
   ```

2. Run the update preparation command:

   ```
   $ openstack overcloud update prepare \
       --templates \
       -r <ROLES DATA FILE> \
       -n <NETWORK DATA FILE> \
       -e <ENVIRONMENT FILE> \
       -e <ENVIRONMENT FILE> \
       ...
   ```

   Include the following options relevant to your environment:

   - Custom configuration environment files (**-e**)

   - If using your own custom roles, include your custom roles (**roles_data**) file (**-r**)

   - If using custom networks, include your composable network (**network_data**) file (**-n**)

3. Wait until the update preparation completes.

3.2. RUNNING THE CONTAINER IMAGE PREPARATION

The overcloud requires the latest OpenStack Platform 15 container images before performing the update. This involves executing the **container_image_prepare** external update process. To execute this process, run the **openstack overcloud external-update run** command against tasks tagged with the **container_image_prepare** tag. These tasks:

- Automatically prepare all container image configuration relevant to your environment.

- Pull the relevant container images to your undercloud, unless you have previously disabled this option.

Procedure

1. Source the **stackrc** file:

   ```
   $ source ~/stackrc
   ```

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

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

3.3. UPDATING ALL CONTROLLER NODES

This process updates all the Controller nodes to the latest OpenStack Platform 15 version. The process involves running the **openstack overcloud update run** command and including the **--nodes Controller** option to restrict operations to the Controller nodes only.

Procedure

1. Source the **stackrc** file:

   ```
   $ source ~/stackrc
   ```

2. Run the update command:

   ```
   $ openstack overcloud update run --nodes Controller
   ```

3. Wait until the Controller node update completes.

3.4. UPDATING ALL COMPUTE NODES

This process updates all Compute nodes to the latest OpenStack Platform 15 version. The process involves running the **openstack overcloud update run** command and including the **--nodes Compute** option to restrict operations to the Compute nodes only.

Procedure

1. Source the **stackrc** file:

   ```
   $ source ~/stackrc
   ```

2. Run the update command:

   ```
   $ openstack overcloud update run --nodes Compute
   ```

3. Wait until the Compute node update completes.

3.5. UPDATING ALL HCI COMPUTE NODES

This process updates the Hyperconverged Infrastructure (HCI) Compute nodes. The process involves:

- Running the **openstack overcloud update run** command and including the **--nodes ComputeHCI** option to restrict operations to the HCI nodes only.

- Running the **openstack overcloud ceph-upgrade run** command to perform an update to a containerized Red Hat Ceph Storage 3 cluster.
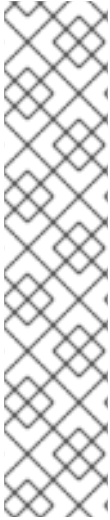
> **NOTE**
>
> Currently, the following combinations of Ansible with **ceph-ansible** are supported:
>
> - **ansible-2.6** with **ceph-ansible-3.2**
>
> - **ansible-2.4** with **ceph-ansible-3.1**
>
> If your environment has **ansible-2.6** with **ceph-ansible-3.1**, run the following commands to update **ceph-ansible** to the newest version:
>
> ```
> # subscription-manager repos --enable=rhel-7-server-rhceph-3-tools-rpms
> # subscription-manager repos --enable=rhel-7-server-ansible-2.6-rpms
> # yum update ceph-ansible
> ```

Procedure

1. Source the **stackrc** file:

   ```
   $ source ~/stackrc
   ```

2. Run the update command:

   ```
   $ openstack overcloud update run --nodes ComputeHCI
   ```

3. Wait until the node update completes.

4. Run the Ceph Storage update command. For example:

   ```
   $ openstack overcloud ceph-upgrade run \
       --templates \
       -e <ENVIRONMENT FILE> \
       -e /home/stack/templates/overcloud_images.yaml
   ```

   Include the following options relevant to your environment:

   - Custom configuration environment files (**-e**)

   - The environment file with your container image locations (**-e**). Note that the update command might display a warning about using the **--container-registry-file**. You can ignore this warning as this option is deprecated in favor of using **-e** for the container image environment file.

   - If applicable, your custom roles (**roles_data**) file (**--roles-file**)

   - If applicable, your composable network (**network_data**) file (**--networks-file**)

5. Wait until the Compute HCI node update completes.

3.6. UPDATING ALL CEPH STORAGE NODES

This process updates the Ceph Storage nodes. The process involves:

- Running the **openstack overcloud update run** command and including the **--nodes CephStorage** option to restrict operations to the Ceph Storage nodes only.

- Running the **openstack overcloud external-update run** command to run **ceph-ansible** as an external process and update the Red Hat Ceph Storage 3 containers.

Procedure

1. Source the **stackrc** file:

   ```
   $ source ~/stackrc
   ```

2. Run the update command:

   ```
   $ openstack overcloud update run --nodes CephStorage
   ```

3. Wait until the node update completes.

4. Run the Ceph Storage container update command:

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

5. Wait until the Ceph Storage container update completes.

3.7. PERFORMING ONLINE DATABASE UPDATES

Some overcloud components require an online upgrade (or migration) of their databases tables. This involves executing the **online_upgrade** external update process. To execute this process, run the **openstack overcloud external-update run** command against tasks tagged with the **online_upgrade** tag. This performs online database updates 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 tagged with the **online_upgrade** tag:

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

3.8. FINALIZING THE UPDATE

The update requires a final step to update the overcloud stack. This ensures the stack's resource structure aligns with a regular deployment of OpenStack Platform 15 and allows you to perform standard **openstack overcloud deploy** functions in the future.

Procedure

1. Source the **stackrc** file:

   ```
   $ source ~/stackrc
   ```

2. Run the update finalization command:

   ```
   $ openstack overcloud update converge \
       --templates \
       -e <ENVIRONMENT FILE> \
       -e <ENVIRONMENT FILE> \
       ...
   ```

   Include the following options relevant to your environment:

   - Custom configuration environment files (**-e**).

   - If applicable, your custom roles (**roles_data**) file (**--roles-file**)

   - If applicable, your composable network (**network_data**) file (**--networks-file**)

3. Wait until the update finalization completes.

CHAPTER 4. REBOOTING THE OVERCLOUD

After performing a minor version update, perform a reboot of your overcloud in case the nodes use a new kernel or new system–level components.

4.1. REBOOTING CONTROLLER AND COMPOSABLE NODES

Complete the following steps to reboot controller nodes and standalone nodes based on composable roles, excluding 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.

5. Check the services. For example:

   a. If the node uses Pacemaker services, check the node has rejoined the cluster:

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

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

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

   c. If the node uses containerized services, check 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.

Procedure

1. Log in to a Ceph MON or Controller node 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
   ```

2. Select the first Ceph Storage node to reboot and log into the node.

3. Reboot the node:

   ```
   $ sudo reboot
   ```

4. Wait until the node boots.

5. Log in to the node and check the cluster status:

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

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

6. Log out of the node, reboot the next node, and check its status. Repeat this process until you have rebooted all Ceph storage nodes.

7. When complete, log into a Ceph MON or Controller node and enable cluster rebalancing again:

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

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

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

4.3. REBOOTING COMPUTE NODES

Complete the following steps to reboot Compute nodes. To ensure minimal downtime of instances in your OpenStack Platform environment, this procedure also includes instructions about migrating instances from the Compute node you want to reboot. This involves the following workflow:

- Decide whether to migrate instances to another Compute node before rebooting the node

- Select and disable the Compute node you want to reboot so that it does not provision new instances

- Migrate the instances to another Compute node

- Reboot the empty Compute node

- 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.

If for some reason you cannot or do not want to 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 will remain down and you must start them manually. 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**. Default value is: 300

For general 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 you want to reboot.

3. From the undercloud, select a Compute Node. Disable the node:

   ```
   $ 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. If you decided not to migrate instances, skip to this step .

6. If you decided to migrate the instances to another Compute node, use one of the following commands:

   a. Migrate the instance to a different host:

      ```
      (overcloud) $ openstack server migrate [instance-id] --live [target-host]--wait
      ```

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

      ```
      (overcloud) $ nova live-migration [instance-id]
      ```

   c. 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.

7. Wait until migration completes.

8. Confirm the migration was successful:

   ```
   (overcloud) $ openstack server list --host [hostname] --all-projects
   ```

9. Continue migrating instances until none remain on the chosen Compute Node.

10. Log in to the Compute Node. Reboot the node:

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

11. Wait until the node boots.

12. Enable the Compute Node again:

```
$ source ~/overcloudrc
(overcloud) $ openstack compute service set [hostname] nova-compute --enable
```

13. Check whether the Compute node is enabled:

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

4.4. REBOOTING HCI COMPUTE NODES

The following procedure reboots Compute hyperconverged infrastructure (HCI) nodes.

Procedure

1. Log in to a Ceph MON or a Controller node and identify the name of the Ceph MON container:

```
$ sudo podman ps | grep -i ceph | grep -i mon

45fe68d340e5  docker-registry.upshift.redhat.com/ceph/rhceph-4.0-rhel8:latest
```

2. Set the **CEPH_MON_CONTAINER** variable to the name of the container:

```
$ CEPH_MON_CONTAINER=ceph-mon-controller-0
```

3. Verify that you can use the **CEPH_MON_CONTAINER** variable to run Ceph commands:

```
$ sudo podman exec $CEPH_MON_CONTAINER ceph -s
```

4. From the Ceph MON or Controller node, disable Ceph Storage cluster rebalancing temporarily:

```
$ sudo podman exec $CEPH_MON_CONTAINER ceph osd set noout
$ sudo podman exec $CEPH_MON_CONTAINER ceph osd set norebalance
```

5. Log in to the undercloud as the **stack** user.

6. List all Compute nodes and their UUIDs:

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

Identify the UUID of the Compute node you aim to reboot.

7. 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
```

8. List all instances on the Compute node:

```
(overcloud) $ openstack server list --host [hostname] --all-projects
```

9. Use one of the following commands to migrate your instances:

    a. Migrate the instance to a specific host of your choice:

    ```
    (overcloud) $ openstack server migrate [instance-id] --live [target-host]--wait
    ```

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

    ```
    (overcloud) $ nova live-migration [instance-id]
    ```

    c. 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.

10. Wait until the migration completes.

11. Confirm that the migration was successful:

```
(overcloud) $ openstack server list --host [hostname] --all-projects
```

12. Continue migrating instances until none remain on the chosen Compute node.

13. Log in to a Ceph MON or a Controller node and check the cluster status:

```
$ sudo podman exec $CEPH_MON_CONTAINER ceph -s
```

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

14. Reboot the Compute HCI node:

```
$ sudo reboot
```

15. Wait until the node boots.

16. Enable the Compute node again:

```
$ source ~/overcloudrc
(overcloud) $ openstack compute service set [hostname] nova-compute --enable
```

17. Verify that the Compute node is enabled:

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

18. Log out of the node, reboot the next node, and check its status. Repeat this process until you have rebooted all Ceph storage nodes.

19. When complete, log in to a Ceph MON or Controller node and enable cluster rebalancing again:

    ```
    $ sudo podman exec $CEPH_MON_CONTAINER ceph osd unset noout
    $ sudo podman exec $CEPH_MON_CONTAINER ceph osd unset norebalance
    ```

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

    ```
    $ sudo podman exec $CEPH_MON_CONTAINER ceph status
    ```